

June 2002

Techniques for Modeling Human Performance in Synthetic Environments: A Supplementary Review

Distribution A:
Approved for public release;
distribution is unlimited.



Authors:

Frank E. Ritter
Nigel R. Shadbolt
David Elliman
Richard M. Young
Fernand Gobet
Gordon D. Baxter

Editor:

Frank E. Ritter

**Techniques for Modeling Human Performance in Synthetic
Environments:**
A Supplementary Review

Techniques for Modeling Human Performance in Synthetic
Environments:
A Supplementary Review

Frank E. Ritter
Nigel R. Shadbolt
David Elliman
Richard M. Young
Fernand Gobet
Gordon D. Baxter

Human Systems Information Analysis Center
Wright-Patterson Air Force Base, Ohio

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 074-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Washington Headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188), Washington, DC 20503

1. AGENCY USE ONLY (Leave blank)		2. REPORT DATE June 2002	3. REPORT TYPE AND DATES COVERED State-of-the-Art Report	
4. TITLE AND SUBTITLE Techniques for Modeling Human Performance in Synthetic Environments: A Supplementary Review			5. FUNDING NUMBERS SPO700-98-D-4001	
6. AUTHOR(S) Frank E. Ritter, Nigel R. Shadbolt, David Elliman, Richard M. Young, Fernand Gobet, & Gordon D. Baxter				
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) HSIAC Program Office AFRL/HEC/HSIAC Bldg 29 2245 Monahan Way Wright-Patterson AFB OH 45433-7008			8. PERFORMING ORGANIZATION HSIAC SOAR 02-02	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES) Defense Technical Information Center ATTN: DOD IAC Program Office (DTIC-AI) 8725 John J. Kingman Road, Suite 0944 Ft. Belvoir, VA 22060-6218			10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES				
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited. Available solely through HSIAC for \$35.00 (US), harcopy, or free-of-charge at http://iac.dtic.mil/hsiac , electronic			12b. DISTRIBUTION CODE A	
13. ABSTRACT (Maximum 200 Words) Selected recent developments and promising directions for improving the quality of models of human performance in synthetic environments are summarized, beginning with the potential uses and goals for behavioral models. The focus is on providing more complete performance, better integration of the models with synthetic environments and with each other (reusability), and improved usability of the models. Within this context, relevant, current work related to modeling is reviewed, such as cognitive modeling of emotion, advanced techniques for testing and building models of behavior, new cognitive architectures including hybrid architectures, and agent and Belief, Desires and Intentions (BDI) architectures. A list of projects with high payoff for modeling human performance in synthetic environments is provided as a conclusion.				
14. SUBJECT TERMS Automated Modeling Support, Cognitive Architectures, Cognitive Engineering, Cognitive Models, Human-computer interaction, Computer Interaction, Intelligent Architectures, Knowledge Acquisition, Models, Synthetic Forces			15. NUMBER OF PAGES 136	
			16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT UNCLASSIFIED	18. SECURITY CLASSIFICATION OF THIS PAGE UNCLASSIFIED	19. SECURITY CLASSIFICATION OF ABSTRACT UNCLASSIFIED	20. LIMITATION OF ABSTRACT	

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18
298-102

Contributors

Gordon D. Baxter, PhD

Department of Psychology
University of York
Heslington
York YO10 5DD
United Kingdom
Email: g.baxter@psych.york.ac.uk
URL: www-users.york.ac.uk/~gdb1

David Elliman, PhD

School of Computer Science & Information Technology
University of Nottingham
Jubilee Campus
Triumph Road
Nottingham NG8 1BB
United Kingdom
Email: dge@cs.nott.ac.uk
URL: www.saruman.cs.notl.ac.uk/index.html

Fernand Gobet, PhD

School of Psychology
University of Nottingham
Nottingham NG7 2RD
United Kingdom
Email: fernand.gobet@nottingham.ac.uk
URL: www.psychology.nottingham.ac.uk/staff/Fernand.Gobet/

Brian Logan, PhD

School of Computer Science & Information Technology
University of Nottingham
Jubilee Campus
Nottingham NG8 1BB
United Kingdom
Email: bsl@cs.nott.ac.uk
URL: www.cs.nott.ac.uk/~bsl/

Frank E. Ritter, PhD

School of Information Sciences and Technology
The Pennsylvania State University
University Park, PA 16801
United States
Email: ritter@ist.psu.edu
URL: www.frankritter.com

Nigel Shadbolt, PhD

Department of Electronics & Computer Science

University of Southampton

Highfield, Southampton 5017 1BJ

United Kingdom

Email: nrs@ecs.soton.ac.uk

URL: www.ecs.soton.ac.uk/~nrs/

Aaron Sloman, PhD

School of Computer Science

The University of Birmingham

Birmingham B15 2TT

United Kingdom

Email: A.Sloman@cs.bham.ac.uk

URL: www.cs.bham.ac.uk/~axs

Richard M. Young, PhD

Department of Psychology

University of Hertfordshire

Hatfield, Hertfordshire AL10 9AB

United Kingdom

Email: r.m.young@herts.ac.uk

URL: www.psy.herts.ac.uk/pub/R.M.Young/

About Human Systems IAC

The Human Systems Information Analysis Center (HSIAC) is the gateway to worldwide sources of up-to-date human-factors information and technologies for designers, engineers, researchers, and human-factors specialists.

HSIAC provides a variety of products and services to government, industry, and academia promoting the use of ergonomics in the design of human-operated and manned systems. HSIAC's primary objective is to acquire, analyze, and disseminate timely information about ergonomics. On a cost-recovery basis, HSIAC will perform the following functions—

- Distribute human-factors and ergonomics technologies and publications.
- Conduct customized bibliographic searches and reviews.
- Prepare state-of-the-art reports and critical reviews.
- Conduct specialized analysis evaluations.
- Organize and/or conduct workshops and conferences.

HSIAC is a Department of Defense Information Analysis Center sponsored by the Defense Technical Information Center, Fort Belvoir, Virginia. It is technically managed by the U.S. Air Force Research Laboratory Human Effectiveness Directorate, Wright-Patterson Air Force Base, Ohio, and operated by Booz Allen Hamilton, McLean, Virginia.

This report, copyright 2003 by Ritter, Shadbolt, Elliman, Young, Gobet, and Baxter.

This material may be reproduced by or for the U.S. Government pursuant to the copyright license under the clause at 52.227-7013 (May 1987).

Additional copies of this state-of-the-art report (SOAR) are available for US \$35 (plus shipping and handling) from—

Human Systems IAC Program Office
AFRL/HEC/HSIAC
2245 Monahan Way, Building 29, Room 206
Wright-Patterson AFB, OH 45433-7008
Tel: (937) 255-4842
Fax: (937) 255-4823
E-mail: sales@wpafb.af.mil
URL: <http://iac.dtic.mil/hsiac>

Table of Contents

Preface.....	xiii
List of Figures.....	xv
List of Tables.....	xv

CHAPTER 1

<u>Tasks and Objectives for Modeling Behavior in Synthetic Environments</u>	<u>1</u>
1.1 The Role of Synthetic Forces	1
1.2 Definition of Terms	2
1.2.1 Synthetic Forces	3
1.2.2 Modular Semi-Automated Forces	3
1.2.3 Frameworks, Theories, Models, and Cognitive Architectures.....	4
1.3 Summary of Modeling Human and Organizational Behavior	5
1.4 What Modeling Human and Organizational Behavior Does Well	5
1.5 Where Modeling Human and Organizational Behavior Can Be Improved	6
1.6 Structure of This Report	6

CHAPTER 2

<u>Current Objective: More Complete Performance.....</u>	<u>9</u>
2.1 Learning	9
2.2 Expertise	10
2.3 Working Memory.....	11
2.4 Emotions and Behavioral Moderators	12
2.4.1 Further Uses of Emotions and Behavioral Moderators	13
2.4.2 Working Within a Cognitive Architecture.....	13
2.4.3 A Sketch of a Computational Theory of Emotions	13
2.5 Errors	15
2.5.1 Training About Errors.....	15
2.5.2 Models That Make Errors.....	16
2.6 Adversarial Problem Solving	16
2.7 Variance in Behavior	17
2.8 Information Overload	18

CHAPTER 3

Current Objective: Better Integration.....19

3.1 Perception	19
3.2 Combining Perception and Problem Solving	20
3.3 Integration of Psychology Theories.....	22
3.4 Integration and Reusability of Models	22
3.5 Summary.....	24

CHAPTER 4

Current Objective: Improved Usability.....25

4.1 Usability of the Models	25
4.2 Desired Accuracy of the Models	26
4.3 Aggregation and Disaggregation of Behaviors	26
4.4 Summary.....	26

CHAPTER 5

Recent Developments for Modeling.....27

5.1 Data Gathering and Analysis Techniques	27
5.2 Advanced AI Approaches.....	28
5.2.1 Genetic Algorithms	28
5.2.2 Tabu Search.....	29
5.2.3 Multiple Criteria Heuristic Search	29
5.3 Psychologically Inspired Architectures	31
5.3.1 Elementary Perceiver and Memoriser	31
5.3.2 Neural Networks	32
5.3.3 Sparse Distributed Memories	33
5.3.4 PSI and Architectures That Include Emotions	33
5.3.5 COGENT.....	34
5.3.6 Hybrid Architectures	35
5.4 Knowledge-Based Systems and Agent Architectures	35
5.5 Architectural Ideas Behind the Sim_Agent Toolkit	37
5.5.1 Cognition and Affect	37

5.5.2 Sim_Agent and CogAff	39
5.5.3 Summary.....	40
5.6 Engineering-Based Architectures and Models.....	41
5.6.1 Apex	41
5.6.2 Simplified Model of Cognition and Contextual Control Model.....	42
5.6.3 Summary.....	43
5.7 Summary of Recent Developments for Modeling Behavior	43

CHAPTER 6

Review of Recent Developments and Objectives: Specific Projects..... 53

6.1 Projects Providing More Complete Performance	53
6.1.1 Gathering Data From Simulations	53
6.1.2 Understanding Expectations of Behavior	54
6.1.3 Including Learning in Models.....	55
6.1.4 Including a Unified Theory of Emotions	55
6.1.5 Including Errors	56
6.1.6 Including a Unified Theory of Personality.....	59
6.1.7 Including a Model of Situation Awareness and Rapid Decision Making.....	60
6.1.8 Using Tabu Search to Model Behavior	60
6.2 Projects Supporting Integration	61
6.2.1 Models of Higher-Level Vision.....	61
6.2.2 Tying Models to Task Environments.....	62
6.2.3 Ongoing Review of Existing Simulations.....	62
6.2.4 Focus on a Flagship Task	63
6.2.5 A Framework for Integrating Models With Simulations	63
6.2.6 A Framework for Integrating Knowledge.....	65
6.2.7 Methods for Comparing Modeling Approaches	65
6.2.8 (Re)Implementing the Battlefield Simulation	66
6.3 Projects Improving Usability.....	68
6.3.1 Defining the Modeling Methodology.....	68
6.3.2 Individual Data Modeling: An Approach For Validating Models	69
6.3.3 Using Genetic Algorithms to Fit Data.....	69
6.3.4 Environments for Model Building and Reuse.....	70
6.3.5 Automatic Model Building.....	71

6.3.6 Improvements to ModSAF	72
6.4 Other Applications of Behavioral Models in Synthetic Environments	72
6.5 Summary of Projects.....	73
APPENDIX A List of Participants	75
APPENDIX B Description of Soar and ACT-R.....	77
B.1 Background of Soar and ACT-R.....	77
B.2 Similarities Between Soar and ACT-R.....	78
B.3 Differences Between Soar and ACT-R	78
B.4 Bibliography for Soar and ACT-R.....	79
Glossary of Acronyms and Abbreviations.....	81
References.....	83
Author Index	103
Subject Index.....	113

Preface

A one-day meeting in March 1999 at Nottingham was convened to explore techniques for modeling human performance in synthetic environments. A list of participants is available as Appendix A. The presentations served as preliminary versions of some chapters of this book. The chapters were expanded based on the day's discussions, extended reflection, and further informal discussion.

Unlike a very similar, earlier review (Elkind, Card, Hochberg, & Huey, 1990) that noted the need to develop theory before applying such models, we are able to conclude that the models presented here are available and useful. The question remaining is how to improve them. We found that the resulting report was usable as a general update to Pew and Mavor's (1998) book, as it reviewed work that was done after their book. In particular, we were able to examine a wider variety of cognitive architectures developed outside the United States. This report also provides a detailed source of further ideas and suggestions for projects. We particularly draw the reader's attention to the importance of the integration and usability of models. Some implications apply more to the United Kingdom and Australia, but nearly all are general.

The report proved popular, so we updated it and looked for a publisher to help disseminate it more widely. Mike McNeese was instrumental in putting us in touch with the Human Systems Information Analysis Center (HSIAC). We are grateful to HSIAC for agreeing to publish this book and preparing it for publication. Comments from Jeffrey A. Landis, HSIAC Publications Manager and Editor, and Dr. Michael Fineberg, HSIAC Chief Scientist, have significantly improved this work. We appreciate their support.

Stephen Croker and Peter Lonsdale provided useful comments and helped assemble these materials. In addition to the workshop participants listed, we thank Angie Barnhill, Tim Barnhill, Christina Bartl, Kevin Gluck, Simon Goss, Ian Greig, Robin Hollands, Nicholas Howden, Jim Jansen, Andrew Lucas, Mike McNeese, Emma Norling, Ralph Rönquist, and Colin Sheppard for their help or comments. Brian Logan and Aaron Sloman, while not listed as authors, did provide material that substantially helped in the preparation of the book. This project was primarily supported by DERA (Bedford, UK) under contract LSA/E20307, and also by DSTO (Australia) and later by the (US) Office of Naval Research (contracts N000140110243, N000140110547, and N000140210021). The conclusions reported here, however, are solely the responsibility of the authors.

Frank E. Ritter
University Park, Pennsylvania
January 2003

List of Figures

Figure 6.1: A functional description of Tac-Air Soar and how it uses ModSAF and a perceptual interface to ModSAF..... 64

List of Tables

Table 1.1. Potential Uses of Models in Synthetic Force Environments2
Table 5.1. Comparison of Architectures45

CHAPTER 1

Tasks and Objectives for Modeling Behavior in Synthetic Environments

There are now numerous models of human behavior in Synthetic Environments (SEs), and they serve a multitude of uses. It is worthwhile considering where and how to improve these models to provide more realistic human behavior. This report provides a more recent review of work following Pew and Mavor (1998), and provides a detailed source of further ideas and suggestions. In addition to noting areas where models could be expanded to include more complete performance, we particularly draw the reader's attention both to the importance of the integration of models (and thus their reuse) and to the usability of models. We will argue that improved usability (and reusability) is necessary for these models to achieve their potential. We extend Pew and Mavor's results by examining architectures (e.g., COGENT, JACK, hybrid architectures) that were not included or available when Pew and Mavor compiled their report, and by summarizing several promising areas for further work that have arisen recently.

This report reflects the biases and specific expertise of the authors as they attempt to identify a wide range of potential problems and provide possible solutions. Some of the proposed projects are high risk and not all of the authors agree that these projects can be accomplished. All agree, however, that if possible, they would be rewarding. Given the diversity of human behavior, there remain many issues not covered here. For example, many aspects of teamwork are important but not examined here. Most of the systems and architectures reported here are continually evolving. Because of the rapid pace of development in this area, our review may underestimate the capabilities of these systems and several of our suggestions may already be incorporated in them.

1.1 The Role of Synthetic Forces

There are several commonly acknowledged uses of cognitive models in synthetic environments. These uses have included at least the range shown in Table 1.1. This is a wide set. Pew and Mavor (1998) focused on the application of synthetic forces to training partly because the major applications and successes of synthetic forces have been in this domain. Further uses of synthetic forces have been outlined in other reviews (Computer Science and Telecommunications Board, 1997; Lucas & Goss, 1999; Synthetic Environments Management Board, 1998).

Table 1.1: Potential Uses of Models in Synthetic Force Environments

- Training leaders
- Joint and combined training
- Training other personnel (e.g., support and logistics)
- Testing existing doctrine
- Testing possible future procurements
- Testing new doctrine
- Serving as a formal, runnable description of doctrine

The user community for synthetic forces would be better served if all these uses were supported by a single system or approach. Currently, the models of behavior in these systems have often been developed without a long-term plan, and are only usable within the simulation for which they were developed. Historically, few single systems have supported more than one or two of the uses noted in Table 1.1. This is wasteful and can lead to different behaviors being taught or used in different simulations when they should be exactly the same behavior. The use of the Distributed Interactive Simulation (DIS) protocol for distributed simulation is a step toward integration, but it does not apply to behavior itself.

While having a single system or approach is highly desirable, there are good reasons why multiple systems are currently used (in addition to a multitude of bad reasons as well). Perhaps the most important reason why there are multiple models of behavior is that existing approaches to modeling cannot support all of the uses in Table 1.1 equally well. Models that focus on aggregate, or large unit behavior, do not support low-level simulations very well. Models that predict average behavior are much less useful for practicing tactics and procedures. Models that are good for training provide detailed data that have to be extensively summarized and aggregated to be of use to planners. Planners and evaluators, for example, may find useful data in large simulations such as the Purple Link exercise, part of STOW97 (further information is available from Ceranowicz, 1998, as well as from www.sticom.amy.mil/STRICOM/DRSTRICOM/DOCATS/), although such simulations cannot yet be convened within an afternoon or even a week to examine how a new platform performs. This report will make suggestions on all of these levels, but it does not intend to be comprehensive.

1.2 Definition of Terms

There are several terms used in this report that have meanings specific to the domain of behavioral modeling. The term model, for example, will refer exclusively to cognitive models, and the term “simulation” will refer exclusively to task simulations. We review these terms here, starting by introducing synthetic forces. Modular Semi-Automated Forces (ModSAF) is briefly explained to provide a common system as a point of reference. We then define the terms we will use with respect to models of behavior.

1.2.1 Synthetic Forces

Synthetic forces exist in military simulations, sometimes alongside real forces that have been instrumented and linked to the simulation. There are now synthetic force simulations covering all of the armed services. Synthetic forces can be separated into two components, physical and behavioral. The physical aspects represent the movement and state of platforms (objects) in the simulation, including such aspects as maximum speed and the set of actions that can be performed in the world. The physical aspects provide constraints on behavior. Simulations of the physical aspects are fairly complete now for most purposes, although they remain important in their own right (Computer Science and Telecommunications Board, 1997; Synthetic Environments Management Board, 1998).

The behavioral aspects of a synthetic force platform determine where, when, and how it performs the physical actions, that is, its behavior. Many human and entity behaviors can be simulated, such as movement and attack, but behavior has been less veridically modeled than physical performance. The next step to increase realism is not only to include further intelligent behavior but also to match more closely the timing and sequence of human behavior when performing the same tasks.

1.2.2 Modular Semi-Automated Forces

Modular Semi-Automated Forces (ModSAF) is a system for simulating entities (platforms) on a simulated battlefield (Loral, 1995). It is perhaps the most widely used behavioral simulator in military synthetic environments. The goal of ModSAF is to replicate the behavior of simulated platforms in sufficient detail to provide useful training and simulation of tactics.

ModSAF includes the ability to simulate the most common types of physical platforms, such as a tank, and external effects on those platforms, like weather and smoke. The terrain is defined in a separate database, which is shared by other simulators in the same exercise using the DIS simulation protocol. Multiple platforms can be simulated by a single ModSAF system.

The local platforms interact with remote platforms by exchanging approximately 20 different types of information packets. Examples of packet types include announcing where the platform is located (the other platforms compute whether the originator can be seen), where radar is being emitted, and where shots are being fired. Thus, the features of the packets vary. Each simulation is responsible for updating its own position and computing what to do with the information in each packet, so that a tank does not directly shoot another tank, for example. Attackers send out projectile packets and the target tank computes that it would be damaged by their projectiles.

Some semi-intelligent behaviors are included in ModSAF through a set of about 20 different simple scripts. These scripts support such activities as moving between two points, hiding, and patrolling.

ModSAF is a large system. It can be compiled into several major versions, including versions to test networks and specific versions for each service. The terrain databases each include up to 1 gigabyte of data. In 1999, simulating multiple entities required a relatively fast workstation (100 MHz+) with a reasonable amount of memory (32 MB+).

A major problem is usability as ModSAF is large and has a complicated syntax. Users report problems learning and using it. A better way to provide its functionality needs to be found or its usability needs to be improved directly.

1.2.3 Frameworks, Theories, Models, and Cognitive Architectures

It is common in cognitive science to differentiate between several levels of theorizing (e.g., Anderson, 1983; 1993, chap. 1) and defining these levels now will help us in the remainder of this report. *Framework* refers to the specification of a few broad principles, with too many details left unspecified to be able to make empirical predictions. For example, the idea that human cognition acts as a production system offers a framework for studying the human mind.

Theory adds more precision to frameworks, and describes data structures and mechanisms that at least allow qualitative predictions to be made. For example, the production system principles presented in Newell and Simon (1972) form a theory of human cognition.

Models are theories implemented as computer programs or represented mathematically to apply to specific situations or types of situations. While generally more limited in their domain of application than theories, models typically provide more accurate, quantitative predictions.

Cognitive architecture has two meanings: (1) specifications of the main modules and mechanisms underlying human cognition, and (2) the computer program implementing these specifications. These meanings are separate and distinct but usually are used as equivalent. Cognitive architectures, as proposed by Newell (1990), offer a platform for developing cognitive models rapidly while keeping the theoretical coherence between these models intact. These cognitive architectures are often seen as equivalent to Unified Theories of Cognition (UTC), a way to pull all that is known about cognition into a single theory. In Appendix B we include brief descriptions of two commonly used cognitive architectures, ACT-R and Soar.

There exists no generally agreed definition of *hybrid architectures*. Some use the term when a cognitive architecture includes symbolic features (e.g., a production system) as well as non-symbolic features (e.g., neural net spreading of activation among memory elements); others, such as Pew and Mavor (1998), use the term when two or more architectures of any kind are combined (e.g., Soar and EPIC). We use the latter definition here because this type of hybrid architecture has become more important and more frequently used.

When comparing theoretical proposals, it is essential to keep in mind the level at which the proposals were formulated. Typically, a framework will cover a large amount of empirical regularities without specifying many details, while a model will cover a small amount of data with great precision. It is generally agreed that models are more useful scientifically than theories or frameworks because they make clear-cut predictions that can be tested with empirical data, and hence, are less amenable to ad hoc explanations (Popper, 1959). Models are, however, harder to create and use.

1.3 Summary of Modeling Human and Organizational Behavior

While the reader is likely to have seen Pew and Mavor's (1998) *Modeling Human and Organizational Behavior*, we briefly review it here to provide background for readers not familiar with it and to provide some useful context. In their book, Pew and Mavor review the state of the art in human-behavior representation as applied to military simulations, with an emphasis on cognitive, team, and organizational behavior. Their book is based on a panel that met for 18 months and drew extensively on a wide range of researchers. It is available as a hardcopy book, as well as online (books.nap.edu/catalog/6173.html).

Pew and Mavor look not just at representing behavior, but also at methods for generating behavior. They provide a review of the uses of models of behavior in synthetic environments. They include a review of the major synthetic environments in use by the U.S. military. These environments are examples of the range of current and potential uses and levels of simulation.

Their book provides a useful summary of integrated (cognitive) architectures. It is comprehensive and clear enough that we have used it to teach undergraduate students. Their summary includes a table comparing the architectures. We will apply the same table to review several additional architectures.

Their book also reviews the important areas to modeling human behavior in synthetic environments. This is a very wide range, encompassing nearly all of human behavior. Their book reviews attention and multi-tasking, memory and learning, human decision making, situation awareness, planning, behavior moderators (such as fatigue and emotions), organizational (small group) behavior, and information warfare (e.g., how the order of information presentation influences decision making). Their book concludes with a framework for developing models of human behavior followed by conclusions and recommendations. Each of these reviews is clearly written and limited only by the space it is allowed. The reviews are quite positive, suggesting that major aspects of behavior are either already being modeled, or can and will be modeled within a few years. This positive tone is in stark contrast to a similar review a decade earlier, which could only note open questions (Elkind, Card, Hochberg, & Huey, 1990).

1.4 What *Modeling Human and Organizational Behavior* Does Well

Pew and Mavor's book is a useful and seminal book for psychology and modeling. Their book is useful because the reviews it provides, while they could be extended, are unusually clear and comprehensive, covering the full range of relevant behavior. It could serve as a useful textbook for professionals in other areas to teach them current results and problems in the areas of psychology and modeling.

Their book is seminal because the authors lay out a complete review of cognition that is widely usable. While their review is similar to Newell's (1990) and Anderson and Lebiere's (1998) reviews, Pew and Mavor's review is not situated within a single architecture; the result is a more global and only slightly less-directed view.

The reviews of the models and data to be modeled together, because of their scope and potential impact, constitute a call to arms for modelers of synthetic forces. The juxtaposition of the data and ways to model them is enticing and exciting. This approach of modeling behavior will significantly influence psychology in general if the modeling work continues to be successful. Models of synthetic forces in the near future will subsume enough general psychology data that they will simply represent the best models in psychology.

1.5 Where *Modeling Human and Organizational Behavior* Can Be Improved

There are surprisingly few problems with Pew and Mavor's review. However, they do not review all of the possible regularities of human behavior. We will add a few additional important regularities and provide further arguments to support many of their main conclusions. They could have referenced, for example, the *Handbook of Perception and Human Performance* (Boff, Kaufman, & Thomas, 1986) and the *Engineering Data Compendium* (Boff & Lincoln, 1988) for a wide-ranging list of existing general regularities in perception and performance (the latter reference has also been put into a CD-Rom version as well, see iac.dtic.mil/hsiac/products/cashe/cashe.html). In the area of human decision making, Dawes' (1994) review is also valuable. Pew and Mavor do not cite a quite relevant report on how this type of modeling is also being developed as entertainment (Computer Science and Telecommunications Board, 1997), and, not surprisingly, they could not report a concurrent similar United Kingdom review (Synthetic Environments Management Board, 1998).

On a high level and early on, they explicitly note that they will not review the usability of behavioral models. We will argue that improved usability is necessary for these models to achieve their potential.

They do not have the space to review all the integrative (cognitive) architectures. While it would be unfair to call this book dated at this point in time, there are already a few architectures worth considering that were not available to them.

They do not dwell on the ability to describe human behavior, instead they focus on how to generate it. There remains some need to be able to describe the behavior before generating it, which we will take up below.

Finally, they did not have the space to lay out very detailed projects to fulfill their short-, medium-, and long-term goals. We provide a more detailed, but still incomplete, set.

1.6 Structure of This Report

Chapter 2 provides amplifications, updates, and additions to Pew and Mavor's list of psychological regularities that should be included in models of human behavior. Chapter 3 notes problems integrating models with simulations as well as problems integrating them with each other to make larger, more complete models. Chapter 4 takes up the issues surrounding usability of behavioral models. Usability of the models themselves was considered to be outside the scope of Pew and Mavor's report (1998, p. 10). We will argue that improving the usability of these models by their creators and other analysts is not only desirable, but necessary for the success of modeling itself. Chapter 5 considers new techniques and cognitive architectures for modeling human behavior in synthetic

environments with respect to the aims of the previous two chapters. Chapter 6 concludes with a list of projects to address problems identified in Chapters 2, 3, and 4 based on the techniques and architectures in Chapter 5.

CHAPTER 2

Current Objective: More Complete Performance

There are a wide range of behaviors that have yet to be incorporated into existing models. Included in this list are numerous additional relevant regularities about human behavior (see Boff & Lincoln, 1988, for a subset). The question that must be addressed is: which behaviors are the most important and most accessible to incorporate? We note here several of the most promising or necessary behaviors to be included next in models of human performance, based on our experiences and previous work.

The suggestions we make later tend to be based on modeling the individual. Much of the behavior being modeled currently in synthetic environments is different because it needs to include small and large groups and is aggregated across time or situations. As smaller time scales and more intricate and fine-grained simulations are developed and used, such as for modeling urban terrorism, the behavioral issues noted here will become more important.

We start with learning. While Pew and Mavor include learning as a useful aspect of performance, we believe learning to be essential. We also expand the case for including models of working memory, perception, emotions and behavioral moderators, and erroneous behavior. We then can examine higher-level aspects of behavior to be considered, starting with integration of models and ending with information overload.

2.1 Learning

Learning is mentioned as important in several ways by Pew and Mavor (1998). Learning (i.e., training) is the largest role of the military in peace time (i.e., rehearsal, p. 30), essential for multi-tasking behavior (pp. 114-115), an important aspect of human behavior (chap. 5), and important within groups (chap. 10). We cover learning again here.

Pew and Mavor mention several of the advantages of learning. There are several additional advantages that we can emphasize. Tactics are influenced by learning. In addition, there is a home-field advantage: working within your own territory, because you know it, makes additional tactics feasible and provides generally improved performance. (Working within your own territory would also provide some additional motivation.)

Including learning in models would provide a mechanism for producing different levels of behavior. Experienced troops, for example, would be different not in some numeric way in that they react faster (although this is probably true), but in a more qualitative way in that they know more and use different strategies. Learning modifies, constrains, and supports the use of computer interfaces (Rieman, Young, & Howes, 1996); similar effects may be found in exploring terrain and implementing tactics in new geographic spaces.

Programming—that is, creating the model directly—may be too difficult. It may be easier for models to learn behaviors than for these behaviors to be programmed directly. This argument has been put forward by connectionist researchers for some time.

Theoretical work in this area of learning has direct implications for training within the military and within schools. Models that learn can be used to understand and optimize learning (Ohlsson, 1992). If we can program models to learn, the behavior and knowledge that result may be different from the initial knowledge that the system started with or from the expert performance that we currently teach. This final knowledge may be useful for teaching. In the case of photocopying (Agre & Shragar, 1990), for example, better strategies arise through practice but are not valuable enough to teach. In military domains, it may be useful to find and then to teach the improved strategies that may arise from grossly extended practice, that is, tactics that are better but that no person has had enough practice to learn before. At that point, explanation of behavior will also become important to understand why the new behavior is useful so that it is trusted.

2.2 Expertise

Expert behavior has an important role to play in models of human performance (Shadbolt & O'Hara, 1997). One of the Western powers' greatest strengths is training in depth *and* breadth. Practice influences speed of processing and error rates, particularly under stress. If synthetic forces are to be used to test doctrine, the effect of training on expertise must be included.

Expert behavior has been studied extensively in recent years and a great deal is known about it (Chipman & Meyrowitz, 1993; Ericsson & Kintsch, 1995; Gobet, 1998; Gobet & Simon, 2000; Hoffman, Crandall, & Shadbolt, 1998). Some essential characteristics of expertise are highly developed perception for the domain material, selective search for solutions in that domain, and a good memory for domain-related material. In most domains, problem-solving behavior (search) differs as well: novices tend to search backward from the situation to find solutions and experts tend to search forward from the situation to find solutions (Larkin, McDermott, Simon, & Simon, 1980). Finally, transfer of expertise to other domains is limited.

Klein and his colleagues (e.g., Klein, 1997) have studied real-time performance in real settings (as opposed to laboratory settings) in detail, and have essentially found that the characteristics mentioned above are also critical in these situations. A number of rather extensive reviews have been undertaken of Klein's approach, which is often referred to as Naturalistic Decision Making (NDM) (e.g., Hoffman & Shadbolt, 1995). A method to elicit this type of knowledge has been developed by Klein and his associates. It is known as the Critical Decision Method and is described in Hoffman et al. (1998). The specifically real-time challenges of acquiring knowledge relating to perceptually cue-rich decision making are discussed in a second Defence Evaluation and Research Agency (DERA), United Kingdom, report by Hoffman and Shadbolt (1996).

Given the fact that it takes a long time to become an expert—the rule of 10 years or 10,000 hours of practice and study is often mentioned (e.g., Simon & Chase, 1973)—the size of the dataset has made it difficult indeed to study real-time learning on the road to expertise. However, real-time learning in simpler problem-solving tasks has been studied and modeling accounts have been provided (Anzai & Simon, 1979; John & Kieras, 1996; Nielsen & Kirsner, 1994; Ritter & Bibby, 2001). Some of these results may apply to expert learning in more complex tasks as well.

While experts vastly outperform non-experts in most domains, exceptions to this rule have been found in domains such as clinical diagnosis, clinical prediction, personnel selection, and actuarial predictions (Dawes, 1988). In these domains, experts perform only slightly better than non-experts, and typically perform worse than simple statistical methods, such as regression analysis. One other aspect of behavior that distinguishes experts from novices is the ability to recover from errors. An important question is to which category military diagnosing and prediction belong because of the uncertainties involved? And, based on this answer, what can be done (either by providing formal tools or by improving training) to remedy this situation and assist error recovery?

The effect of learning local environments and strategies (own and opponent's) must also be included. Having learned the local terrain probably explains much of the home-field advantage. How does this learning occur?

Within the sub-field of knowledge-engineering there have been considerable efforts to produce methodologies for the acquisition, modeling and implementation of knowledge-intensive tasks. It is a moot point whether the resulting decision-support systems are cognitively plausible. Nevertheless, these methodologies now provide powerful ways of constructing complex systems that exhibit task-oriented behavior. To this end, anyone engaged in engineering large-scale synthetic environments should look at the principles laid down in the most recent of this work. The most accessible source is probably Schreiber et al. (2000).

2.3 Working Memory

Central to all questions about human cognition and performance is the role of working memory. Working memory is implicated in almost all aspects of cognitive performance (Boff & Lincoln, 1986, Sec. 7; Just & Carpenter, 1992; Newell & Simon, 1972; Wickens, 1992). It is widely agreed that limitations of working memory are a major determinant of limitations of cognitive performance. Definitions of working memory are varied but for present purposes we can take it to refer to the mechanisms that maintain and provide access to information created or retrieved during the performance of a task.

Modern approaches to the psychological study of human working memory often take as their starting point the famous paper by Miller (1956) and argue that people can retain only around "7 +/- 2" items in short-term memory. Later work has tended to revise that estimate downwards, towards three to four items of unrelated information (Crowder, 1976; Simon, 1974).

A more recent and influential line of work by Baddeley (1986, 1997) presents working memory as a dual system for the rehearsal of information, consisting of (1) a phonological loop, that contains approximately 2 seconds of verbalizations, for the rehearsal of phonological, acoustic, or articulatory information (e.g., useful for repeating a phone number until you dial it); and (2) a visual-spatial scratchpad, with a smaller and less-determined capacity (e.g., useful when searching for an object that you have just seen), to play an analogous role for the maintenance of pictorial and spatial information.

Other approaches within experimental psychology place more emphasis on the role of working memory in both storing and manipulating temporary information (Daneman & Carpenter, 1980; Just & Carpenter, 1992). An important recent extension to the notion of working memory comes from the study of expertise, where Ericsson and Kintsch (1995) argue that after extensive practice in a particular domain people can, through specialized *retrieval structures*, use long-term memory for the rapid storage of temporary information (i.e., long-term working memory).

A recent book (Miyake & Shah, 1999) reviews a range of current approaches to the modeling of working memory, although many of the models do not have the explicitness and generality needed to support the simulation of human performance in complex tasks. Of those that do, their view of working memory varies widely. Some, such as ACT-R (Anderson & Lebiere, 1998) and CAPS (Just & Carpenter, 1992), consider working memory not as a separate structural entity but rather as an activated region of a larger, more general memory system, in which the limitations of working memory derive from a limited total quantity of activation. Just and Carpenter (1992), and more recently ACT-R models, have extended that view to the modeling of individual differences in working memory where different people are assumed to have different maximum quantities of available activation (Daily, Lovett, & Reder, 2001; Lovett, Daily, & Reder, 2000). A number of these ideas are put together by Byrne and Bovair (1997) who modeled (in CAPS) the way that a class of performance errors, in which people forget to complete subsidiary aspects of a task (such as removing the original from a photocopier), is affected by working memory load.

In contrast to these resource-limited models, Soar (Laird, Newell, & Rosenbloom, 1987; Newell, 1990) imposes no structural limitation on working memory. Using Soar, Young and Lewis (1999) explore the possibilities of working memory being constrained not by physical resources but by functional limitations and by specific kinds of similarity-based interference.

In summary, the current position is that human performance is known to be highly dependent on working memory and working memory load, and to be susceptible to factors such as individual differences (Just & Carpenter, 1992), distractions (Byrne & Bovair, 1997), emotion and stress (Boff & Lincoln, 1988), and expertise (Ericsson & Kintsch, 1995). Many existing models of human performance (e.g., as reviewed in Pew & Mavor, 1998) do not directly model the role of working memory. Models exist (Miyake & Shah, 1999), and some approaches to cognitive modeling (ACT-R, CAPS, Soar) have potential for improving predictions of human performance in realistic task situations by including more accurate theories of memory. There remains a need for the investigation and development of more explicit and complete models, with broader scope, of the role of working memory in human performance.

2.4 Emotions and Behavioral Moderators

Emotion, affect, motivation, and other behavioral moderators are increasingly being seen as factors that can and often do influence cognition. This view has received attention among a range of computer scientists and psychologists. Pew and Mavor (1998, chap. 9) lay out an initial case for including emotion as an internal moderator of behavior. The British HCI Group sponsored a one-day meeting on “Affective Computing: The Role of Emotion in

Human Computer Interaction” that attracted 70 people to University College, London (Monk, Sasse, & Crerar, 1999). Picard’s (1997) recent book provides a useful review of emotions and computation in general. Sloman’s (1999) review of the book and Picard’s (1999) response are useful summaries. A further case is also made in the section on the Sim_Agent Toolkit. We present here an additional argument for including a model of emotions and behavioral moderators in models of synthetic forces, note two potential problems with existing models, and sketch an initial theory.

2.4.1 Further Uses of Emotions and Behavioral Moderators

Models of emotions and behavioral moderators may be necessary for modeling non-doctrinal performance such as insubordination, fatigue, errors, and mistakes. Many authors have also noted the role of emotion in fast, reactive systems (Picard, 1997, provides a useful overview). Individual differences in emotions may be related to personality and differences in problem solving. That is, the range of emotions may be best explained as an interaction that arises between task performance and situation assessment and an agent’s likes, desires, and personal cognitive style. An argument is starting to be put forward that changes in motivation based on temporally local measures of success and failure may help problem solving (Belavkin, 2001; Belavkin & Ritter, 2000; Belavkin, Ritter, & Elliman, 1999).

2.4.2 Working Within a Cognitive Architecture

Emotions arise from structures related to cognition and should be closely related to and based on cognitive structures. All of the arguments for creating a unified theory of cognition (Anderson, Matessa, & Lebiere, 1998; Newell, 1990) also apply to creating a unified theory of emotion as well. The effects of emotions and other behavioral moderators on cognition are presumably not task-specific, so their implementation belongs in the architecture, not in the task knowledge.

Theories of emotions should thus be implemented within a cognitive architecture. This will allow them to realize all the advantages of being within a cognitive architecture, including being reusable and being compared to and incorporated within other models. Some models of emotions have been built within a cognitive architecture (Bartl & Dörner, 1998; Belavkin, Ritter, & Elliman, 1999; Franceschini, McBride, & Sheldon, 2001; Gratch & Marsella, 2001; R. Jones, 1998; Rosenbloom, 1998). Being created within an information-processing model has required them to be more specified than previous theories. Being part of a model that performs the task has also allowed them to make more predictions.

2.4.3 A Sketch of a Computational Theory of Emotions

An important aspect of cognition is to process sensory information, assign meaning to it, and then decide upon a plan of action in response. This is a real-time process in which new sensory information arrives continuously. This view is similar to the view put forward by Agre and Chapman (1987) about representationless thinking. The plan must therefore be dynamically reconfigurable and will often be abandoned in favor of a better plan midway through its execution. Elliman has a speculative view of the role of emotions in cognition,

similar to Rasmussen's (1998) stepladder framework of behavior, which makes the following assumptions:

- The amount of sensory data available at any moment is too large for attention to be given to more than a small fraction of the data.
- The conscious consideration of the results of perception is an expensive process in terms of the load on neural hardware and also time-consuming.
- Most sensory processing is unconscious in its early stages in order that expensive conscious processes need consider only the *results* of perception. These results might include labeled objects with a position in space, for example “a tank moving its turret in that clump of trees.” Conscious processes might well add further detail such as the type of tank and the range of its gun.
- Attentional mechanisms are needed to direct the limited high-level processing to the most *interesting* objects. These may be novel, brightly colored, fast-moving, or potentially threatening.
- Planning is an especially heavy computational process for the human mind and one that is difficult to carry out effectively under combat conditions. (Perhaps the best way to explain why military doctrine is useful is that it distills the best generic practice and trains the soldier to behave in a way that might well have been a chosen and planned behavior if the individual had the time and skill to formulate the action himself. The danger is that no doctrine can envisage all scenarios in advance and, on occasion, the use of doctrine in a rigid manner may be harmful.)
- From an evolutionary perspective this system of unconscious processing of sensory input, attentional mechanisms, and cognitive planning (together with speech-based communication) is a masterstroke of competence for survival. However, it has one crippling disadvantage—it is too slow to react to immediate and sudden attack.

Rapid reaction to possible threat without the time for much cognitive processing is clearly of huge value. In this framework emotion can be seen as kind of labeling process for sensory input. Fear particularly fits this pattern and is a label that causes selected sensory input to literally *scream* for attention. For this process to work rapidly it needs to be hardwired differently than higher-level cognitive processes. There is strong evidence that the amygdala is intimately involved in the perception of threat and able to trigger the familiar sensation of fear (e.g., Whalen, 1999). If this organ of the brain is damaged, individuals may find everyday events terrifying while not perceiving any need for alarm in life-threatening situations.

This rapid, emotive response to sensory data is relatively crude and prone to false alarms. Reactive behavior is triggered that may be involuntary, for example, the startle reaction and physiological changes due to the release of noradrenalin. After the reaction response, it takes time for cognitive processes to catch up and make a more informed assessment of the situation and actual threat. If this emotive, reactive stimulation is excited in a chronic manner then susceptible individuals may become less effective, with impaired ability to think and plan clearly. Any kind of anxiety is a form of stress. Because individuals have a finite capacity for absorbing it, excessive stress results in fatigue.

2.5 Errors

Ideally, military behavior is normative, that is, what *is* done is what *should* have been done. Human behavior does not always match the normative ideal of military behaviors. One of the most important aspects of human performance, which has often been overlooked in models of behavior and problem solving, is errors (although see, for example, Cacciabue, Decortis, Drozdowicz, Masson, & Nordvik, 1992; Freed & Remington, 2000; Freed, Shafto, & Remington, 1998). There is a consensus building about the definition of errors—for most people an error is something done that was not intended by the actor, that was not desired, and that placed the task/system beyond acceptable limits (e.g., Senders & Moray, 1991).

Part of the reason for omitting errors from models of behavior is the fallacy that they are produced by some special error-generating mechanism that can be bolted on to models once they are producing correct behavior on the task at hand. Often, however, the actions that precede errors would have been judged to be correct if the circumstances had been slightly different. In other words, as Mach (1905/1976) observed, knowledge and error both stem from the same source.

Evidence shows that novices and experienced personnel will often make the same errors when exposed to the same circumstances. The difference lies in the ability to notice and recover from these errors. Experienced personnel are more successful at mitigating errors before the full consequences arise. In other words, it is the management of errors that is important and needs to be trained (Frese & Altmann, 1989), rather than vainly trying to teach people how to prevent the inevitable.

2.5.1 Training About Errors

In any complex, dynamic environment, such as a military battlefield, the consequences of uncorrected errors are potentially disastrous. While normally a string of mistakes is required to create a disaster, the rapid pace of the battlefield and adversaries allows single mistakes to become more catastrophic.

There is, therefore, a real need to learn how to manage errors in an environment in which the consequences are less severe. An advantage of using synthetic environments is that comparative novices can experiment in unfamiliar situations, with restrictions approximating the real environment in time, space, enemy capabilities, and so on, but with the knowledge that the consequences of any errors can be recovered. In addition, multiple scenarios can be played out over a compressed time period, thereby providing the novice with a variety of experiences that would take many years to accumulate through exposure to situations in the real world. This can be a great training aid, literally giving years of experience in far less time. When novices were trained in aircraft electrical-system troubleshooting using a simulated system, they were able to acquire years of experience in months because the tutor let them practice just their diagnostic skills without practicing their disassembly skills (Lesgold, Lajoie, Bunzon, & Eggan, 1992).

2.5.2 Models That Make Errors

There are several process models complete enough to make errors, depending to some degree on the definition of error. Models that include errorful behavior exist in EPAM (Feigenbaum & Simon, 1984; Gobet & Simon, 2000), ACT-R (Anderson, Farrell, & Sauer, 1984; Anderson & Lebiere, 1998; Lebiere, Anderson, & Reder, 1994) and Soar (Bass, Baxter, & Ritter, 1995; Howes & Young, 1996; Miller & Laird, 1996), although each generates errors in different ways and at different levels. Fewer models exist that model error recovery, although this is clearly the next aspect to model.

A problem with models and humans is that the erroneous behavior is often task-specific; given a new task, both models and humans might not generate the same behavior. In other words, the erroneous behavior arises as a result of the combination of human, technological, and organizational (environmental) factors. Vicente (1998) delineates some of the problems in this area.

There are various taxonomies of errors that could be incorporated into models of performance. There are also other constraints that reduce the level of performance that are worth exploring, including working memory (Young & Lewis, 1999), attention, and processing speed due to expertise.

2.6 Adversarial Problem Solving

Adversarial problem solving is different from simple problem solving and makes additional requirements for modeling behavior in synthetic environments. Planning is not done within a static environment, but done in an environment with active adversaries.

Research on adversarial problem solving (e.g., Chase & Simon, 1973; de Groot 1946/1978; Gobet & Simon, 2001; Newell & Simon, 1972) has identified several aspects of cognitive behavior that have been shown to generalize to other domains, including the military domain (Charness, 1992). A key result is that players do not follow a strategy such as *minimax* but that they satisfice (Simon, 1955), that is, they satisfy themselves with a good-enough solution, which can be far from the optimal solution (de Groot & Gobet, 1996; Gobet & Simon, 1996a). This satisficing behavior can be explained by the processing and capacity limits of human cognition, such as the time to learn a new chunk or the capacity of short-term memory (Newell & Simon, 1972).

A second, related aspect is that a player's search is highly selective: only a few branches of the search tree are explored. The choice of subspace to search seems to be constrained by pattern-recognition mechanisms (Chase & Simon, 1973; Gobet, 1998; Gobet & Simon, 1996a). A consequence is that misleading perceptual cues may result in the exploration of an incorrect subspace. For example, Saariluoma (1990) reported that chess masters found a suboptimal solution when the features of the position led them to look for a standard, although inferior, subspace. The consequence for understanding combatant behavior is that pattern recognition may influence the course of action chosen as much as the detail of the way the search is carried out. In fact, de Groot (1946/1978) did not find differences in the macrostructure of search of chess players at different skill levels.

A third important result is that chess players re-investigate the same sequence of actions several times, interrupted or not by the analysis of other sets of actions. De Groot (1946) has

called this phenomenon *progressive deepening*. It is related to the selective search shown by experts in other areas (Charness, 1991; Ericsson & Kintsch, 1995; Gobet & Simon, 1996a; Hoffman, 1992). De Groot and Gobet (1996) propose that progressive deepening is due both to the limits of human cognition (limited capacity of short-term memory, slow encoding time in long-term memory) and that with this searching behavior, information gathered at various points of the search may be propagated to other points, including previously visited points (this could not be done with a search behavior such as *minimax*).

These features of cognition, identified in adversarial problem solving, also occur in Rapid Decision Making (RDM) in domains such as firefighting, combat, and chess players in time-trouble. Interestingly, the model developed by Klein and his colleagues (see Klein, 1997, for a review) singles out the same features as the model developed by Chase and Simon (1973) to explain expert chess-playing: pattern recognition, selective search, and satisficing behavior.

While some aspects of adversarial problem solving are well understood, others have yet to be studied in any depth. Such aspects include the way the function used to evaluate the goodness of a state (the evaluation function) changes as a function of time, the link between the evaluation function and pattern recognition, or the learning of domain-specific heuristics, which all have direct implications for combat behavior.

Relatively little research has been done on how players take advantage of the thinking particularities of their opponent, in particular, by trying to outguess him or her. Jansen (1992) offers interesting results. He has developed a computer program that takes advantage of some features and heuristics of human cognition in simple chess endgames, such as the tendency, in human players' search, to avoid moves that lead to positions with a high-branching factor, and to prefer moves that lead to forced replies. Using these features and incorporating them in its evaluation function, the program was able to win faster (in won positions) or to avoid defeat (in lost positions) more often against human players than by using a standard alpha-beta search. In principle, such an approach could be extended to include both skill-related and individual differences in synthetic environments.

In comparison to perception and memory in games, relatively little computer modeling of human behavior has been done with adversarial problem solving (if one excludes pure Artificial Intelligence [AI] research, in which adversarial problem solving has been a favorite subject of research). One may mention the previous work of Simon and colleagues (Baylor & Simon, 1966; Newell, Shaw, & Simon, 1958), and the programs of Pitrat (1977), Wilkins (1980), and Gobet and Jansen (1994). All of these programs were created for chess and most cover only a subset of the game.

There are implications of adversarial search variation for performance (i.e., how well a planner models an opponent). This would be a natural place to model various levels of experience in opponents.

2.7 Variance in Behavior

Including more variety in how a model performs a task is one of the next steps for improving the realism of synthetic forces. Currently, many models will execute a task the same way every time and for every equivalent agent. In the real world, this is not the case. The choice of strategies and the ordering of substrategies will vary across agents and vary

for a given agent across time. This lack of variance makes adversaries and allies too predictable in that they always do the same thing.

Including variance in behavior is also necessary when behavior is less predictable. Novices, with less knowledge, have greater variance in behavior (Rauterberg, 1993). In the past, variance was intentionally suppressed in simulations because it was thought that variance in real behavior was suppressed through doctrine and training. Accounting for variety in behavior is of increasing importance when modeling less-prepared and less-trained forces, and now for improving model accuracy as variance in real behavior is admitted.

Variance in behavior is also important when modeling non-combatant agents, such as white forces and civilians. These agents may be producing their behaviors deterministically, but the determiners are often hidden from other agents, making them appear relatively unpredictable. Finally, the ability to model a variety of behaviors is necessary for sensitivity analysis.

Variance will arise out of several factors. It may arise from different levels of expertise, which is covered above. It may arise from different strategies, which will require including multiple strategies and noting where orders are less likely to be followed and when panic results in orders being ignored. Variance may also arise as a type of error, such as applying a right action in the wrong circumstances.

In any case, variance in agent behavior in synthetic environments particularly needs to be included in training materials. Humans are very good pattern-recognizers—although they do not always look for or know the right pattern—and will take advantage of models that do not vary their behavior. The real opponents may not be so predictable.

2.8 Information Overload

Problems with information overload have been noted numerous times (e.g., Woods, Patterson, Roth, & Christoffersen, 1999). Hoffman and Shadbolt (1996) provide a review of work on information overload in real-time, high-workload military contexts. They also discuss challenges that information overload raises for knowledge acquisition in the context of synthetic forces environments.

Problems resolving clutter, workload bottlenecks, and finding significance in incoming data, are not yet problems for many models of human performance. Currently, most cognitive and synthetic force models do not face information overload. The situation has more typically been of a model seeing only a limited set of information and knowing how to perform only one or a few tasks.

In the near future, the models will have more complex simulated eyes as well as more knowledge to interpret the eyes' input. This will lead to more incoming information with a more difficult problem of deciding which objective to pursue next and how to choose the best strategy based on a larger set of knowledge and perceptual inputs. We will also find that models will start to have trouble with information overload, clutter, and situation assessment. Their tactics in this area will be particularly important when there are time pressures, which are common in synthetic environments and the worlds they model.

CHAPTER 3

Current Objective: Better Integration

There are theoretical and practical problems integrating models with simulations and with other models. The problems can appear to be simply software issues, but deeper theoretical issues often go hand-in-hand with these problems. We thus note a few of these problems in getting models to interact with simulations as well the basic problem of aggregating models.

3.1 Perception

At least since de Groot's early work (1946), perception has been deemed to play an essential role in cognition. Neisser (1976, p. 9) aptly summarizes it as "perception is where cognition and reality meet." This point of view has been buttressed in recent years with the emphasis given by Nouvelle AI (e.g., Brooks, 1992), which is based on reactive architectures, perceptual mechanisms, and on their coupling with motor behavior. Neuroscience (e.g., Kosslyn & Koenig, 1992) teaches that, due to evolutionary pressure, a large part of the brain deals with perception (mainly vision); hence, an understanding of perception is essential for understanding the behavior of combatants.

Perception-based behavior offers a series of advantages: it is fast, attuned to the environment, and optimized with respect to its coupling with motor behavior. However, its disadvantages include its tendency to be stereotyped and to lack generalization. In addition, from the point of view of the modeler, it is a difficult behavior to simulate well. This is in part due to the fact that low-level perception is still poorly understood (Kosslyn & Koenig, 1992), although recent progress in robotics and agent behavior give examples of successful implementation of basic perceptual mechanisms for use by cognition (e.g., Brooks, 1992; Zettlemoyer & St. Amant, 1999; and St. Amant & Riedl, 2001).

Perception may be seen as the common ground where various aspects of cognition meet, including motor behavior, concept formation and categorization, problem solving, memory, and emotions. In several of these domains, computer simulations illustrating the role of perception have been developed.

Brooks (1992) and others have investigated the role of perception in motor behavior with simple insect-like robots. The link between concept formation and (high-level) perception has been studied using the EPAM architecture (Gobet, Richman, Staszewski, & Simon, 1997). The role of perception in problem solving has been studied using Chunk Hierarchy and REtrieval Structures (CHREST), a variation of EPAM (Gobet, 1997; Gobet & Jansen, 1994) that also accounts for multiple memory regularities. Eye movements are simulated in detail in CHREST but not the low-level aspect of perception. (We will deal with the relation between problem solving and perception in Sec. 3.2.)

A more detailed simulation of low-level aspects of perception, such as feature extraction, is an important goal for the future of research on the relation of perception to other aspects of cognition. In addition, little work has been done on modeling perception in dynamically changing environments and on the effects of stress, emotion, motivation, and group factors on perception.

It is useful to separate *perception* from *cognition* in modeling human performance. The border between the model of the person and their environment can (arguably) be drawn at the boundary between cognition and perception, with perception belonging to a large extent in the environment model. This may be true for psychological reasons (Pylyshyn, 1999). It is also true to support tying models to simulations and for use of the resulting knowledge by cognition in problem solving (Ritter, Baxter, Jones, & Young, 2000). The typical acts performed by perception and motor action, such as determining the objects in view, their shapes and sizes, and then manipulating them, are most easily performed where the objects reside. This forces the implementation of theories of interaction into the simulation language instead of the modeling language.

It would be useful to have realistic stochastic distributions of differences in perception among individual agents, and also the ability to augment perception with instruments from field glasses to night sights. These devices could be modeled as *plug-ins* to the perception model. Models of perception in synthetic environments are typically simple, being a function of distance from observer to object (e.g., if there is a clear line of sight and the absence of cover and smoke). On the other hand, human vision changes in important ways with the ambient level of light and with the part of the retina on which an image falls. The edges of the retina are particularly sensitive to the detection of a moving object, while the fovea has the best resolution for identifying distant objects and is most sensitive to color. The distance at which an object can be seen depends on its brightness, its size, and its contrast to the background as well as the permeability of the air to light. Thus, a detonation will be visible from a much greater range than a moving tank, which in turn will be much easier to spot than a motionless, camouflaged soldier.

Situation awareness is a term that is still the subject of much debate in the human factors and ergonomics communities (e.g., see the Special Issue of *Human Factors*, Volume 37, Issue 1). Pew and Mavor (1998) consider situation awareness to be a key concept in the understanding of military behavior. We agree, but also believe that situation awareness should be modeled at a finer level of detail than is currently often done (see Pew & Mavor, 1998, chap. 7, for a current review).

3.2 Combining Perception and Problem Solving

Pew and Mavor (1998) note that an important constraint on problem solving is perception, but do not explore this in detail. As mentioned in our discussion on expertise, perception plays an important role in skilled behavior—experts sometimes literally *see* the solution to a problem (de Groot, 1946/1978).

We may use Kosslyn and Koenig's (1992) definition: higher-level visual processing involves using previously stored information; lower-level visual processing does not involve such stored information and is driven only by the information impinging on the retina. We

focus here on higher-level perception and, thus, we will not consider mechanisms used for finding edges, computing depth, and so on.

Neisser's *Cognition and Reality* (1976) describes what is often referred to as the perceptual cycle. This approach underpins a vast amount of the cognitive engineering literature and research. At its simplest, the perceptual cycle is a cycle between the exploration of reality and representing this reality as schemas (in the general sense). Schemas direct exploration (perceptual, haptic, etc.) that involves sampling the object (looking at the real world), which may alter the object, which means that the schemas have to be modified. (See Neisser, 1976, p. 21, or p. 112 for a more complete description.) This work suggests that an important aspect of behavior has been missing from many theories and models of problem solving that have not included perception.

It is natural that researchers have attempted in recent years to combine perception and problem solving in artificial systems. One can single out three main approaches: robotics, problem-solving architectures incorporating perception, and perceptual architectures being extended to problem solving.

In robotics, Nouvelle AI has attempted to build robots able to carry simple problem-solving behavior without explicit planning by linking sensor and motor abilities tightly (e.g., the behavior-based architecture of Brooks, 1992). Robots based on this approach are excellent at obstacle-avoiding behavior. It is, however, unclear how far this approach can be extended to more complex problem solving without incorporating some sort of planning.

Including perception in behavioral models is a useful way to add natural competencies and limitations to behavior. Pew and Mavor note that there are few models of how perception influences problem solving. Their summary can be extended and revised in this area, however. We have seen in existing cognitive models (Byrne, 2001; Chong, 2001; de Groot & Gobet, 1996; Gobet, 1997; Jones, Ritter, & Wood, 2000; Ritter & Bibby, 2001; Salvucci, 2001) and in AI models (Elliman, 1989; Grimes, Picton, & Elliman, 1996; St. Amant & Riedl, 2001) that perception is linked to and can provide behavioral competencies and restrictions on problem solving. While Pew and Mavor note that they are unaware of any attempt in Soar to model the detailed visual perceptual processes in instrument scanning (Pew & Mavor, 1998, p. 181), such models exist (Aasman, 1995; Aasman & Michon, 1992; Bass et al., 1995), and some are even cited by Pew and Mavor (1998, p. 95) for other reasons.

The Soar model reported by Bass et al. (1995) scans a simple air-traffic control display to find wind velocity. The model learns (chunks) this information and uses it and the display to track and land a plane through airport air traffic control. The model then reflects on what it did to consider a better course of action. This model shows tentative steps towards using Soar's learning mechanism for situation learning and assessment based on information acquired through active perception (see Pew & Mavor, 1998, p. 197). Modeling visual cognition within Soar is ongoing at the University of Southern California's Information Sciences Institute (USC/ISI; Hill, 1999) and at the Pennsylvania State University.

The EPAM architecture (Feigenbaum & Simon, 1984), the initial goal of which was to model memory and perception, has recently been extended into a running production system (Gobet & Jansen, 1994; Lane, Cheng, & Gobet, 1999). The chunks learned while interacting

with the task environment can later be used as conditions of productions. The same chunks are also used for the creation of schemas and for directing eye movements.

Recently, there have been several attempts to move the perception component from models into the architectures, regularizing and generalizing the results in the process. Prominent cognitive architectures Soar and ACT-R have been extended to incorporate perceptual modules, and PSI also has a perceptual module. With Soar, a perceptual module is available based on EPIC (Chong & Laird, 1997) and another based loosely on a spotlight theory of attention (Ritter et al., 2000). With ACT-R, two perceptual modules have been developed independently: the Nottingham architecture (Ritter et al., 2000) and ACT-R/PM (based on but also extending EPIC; Byrne, 1997, 2001). This approach creates situated models of cognition, that is, models that interact with (simulations of) the real world.

None of these approaches has been tested with complex, natural, and dynamically changing environments. The robotics approach is the only one currently demonstrated to cope with natural, albeit rather simple, environments. The two other approaches can interact with computer interfaces that are complex and dynamic (e.g., Salvucci, 2001).

3.3 Integration of Psychology Theories

A glance at almost any psychology textbook reveals that the study of human cognition is conventionally divided into topics that are presented as if they have little to do with each other. There will be separate chapters on attention, memory, problem solving, and so on. However, the range and variety of tasks undertaken by people at work, and also those tackled by synthetic agents, typically require the application and interplay of many different aspects of cognition simultaneously or in close succession. Interacting with a piece of electronic equipment, for example, can draw upon an agent's capacity for perception, memory, learning, problem solving, motor control, decision making, and many more capabilities. The question of how to integrate these different facets of cognition is therefore an important one for the simulation of human behavior.

Integrating theories across different topics of cognition is an issue that has rarely been addressed directly and provides an important focus for future work. Agents in synthetic environments (e.g., R. Jones, Laird, Nielsen, Coulter, Kenny, & Koss, 1999) implicitly integrate multiple aspects of behavior. What research exists has been carried out, appropriately enough, under the heading of unified theories of cognition using architectures such as Soar and ACT-R. Soar offers a promising basis for such integration. Its impasse-driven organization enables it to access different areas of cognitive skill as the need arises, and its learning mechanism (which depends on cognitive processing in those impasses) enables relevant information from the different areas to be integrated into directly applicable knowledge for future use. ACT-R also integrates multiple components.

3.4 Integration and Reusability of Models

Integration of theories can be also viewed as integration of models as software, sometimes called *reuse*. It has been true for years that reuse is important; this is true for two fundamental reasons. First, reuse saves effort. In the field of object-oriented software development, figures are often quoted for the costs associated with development with reuse

in mind. The extra time spent in initial development is something like 20%. When the code is reused, an application can be created in 40% of the development time for new code. Second, and perhaps more importantly in these domains, reuse ensures consistency across simulations and time, particularly important when creating unified theories of cognition.

There are also serious problems restricting the reuse of cognitive models. Cognitive models are not generally reused, even when they have been created in a cognitive architecture designed to facilitate their reuse. There are exceptions. Pearson's Version 2 of his Symbolic Concept Acquisition model and its explanatory displays is an exception (available at ai.eecs.umich.edu/soar/soar-group.html). Other exceptions include PDP toolkits such as O'Reilly's PDP++ (www.cnbc.cmu.edu/PDP++/PDP++.html). But, overall, cognitive modeling does not have the level of system reuse and visual displays that the AI and expert systems communities now take for granted. This problem is being noticed by others as well (Wray, 2001).

There are some examples of reuse that should be emulated and expanded. ACT-R now maintains a library of existing models (act.psy.cmu.edu). We have found that the mere existence of a library of student models (www.nottingham.ac.uk/pub/soar/nottingham/) has led to increasingly better student projects. Work by Young (1999) on building a zoo of runnable cognitive models is another example of such use done broadly. There is little reason to believe that these results would not scale up. These improvements to the modeling environment have helped move learning Soar (Ritter & Young, 1999) and ACT-R (Anderson & Lebiere, 1998) from being a lengthy apprenticeship to being something that can be taught in undergraduate courses.

Such integration is illustrated most clearly in a model of natural language sentence processing (Lewis, 1993), in which lexical, syntactic, semantic, pragmatic, and domain-specific knowledge are brought together in learned rules (Soar chunks) to guide language comprehension. Probably the model that has gone furthest in demonstrating this kind of integration is the cognitive model of the NASA Test Director, the person responsible for coordinating the preparation and launch of the space shuttle. Nelson, Lehman, and John (1994) describe a Soar model of a fragment of the Test Director's performance, which incorporates problem solving, listening to audio communications, understanding language, speaking, visual scanning (through a procedure manual), page turning, and more. Such integrated models are also starting to be created in ACT-R (Anderson & Lebiere, 1998).

Integration of a slightly different flavor—across capabilities rather than across textbook-like topics of cognition—is illustrated in another Soar model, this one being of exploratory learning of an interactive device (Rieman et al., 1996). At first glance, it might seem that exploratory learning is not especially relevant to the human behavior that is, apart from questions of training, the main focus of this report. Fighter pilots and tank commanders are highly trained and expert individuals, and presumably do not learn significantly from further experiences. However, component skills such as comprehending a novel situation, looking around to discover relevant options, and assessing a course of action—which are fundamental components of expert skill—are also precisely what are required for exploratory learning and reactive planning in uncertain environments.

Rieman et al. (1996) describe the IDXL model, which models an experienced computer-user employing exploratory learning to discover how to perform specified tasks with an unfamiliar software application. IDXL searches both the external space provided by the

software and the internal space of potentially relevant knowledge. It seeks to comprehend what it finds and approximates the *rationally optimal* strategy (Anderson, 1990) for exploratory search. A typical sequence of interrelated capabilities would be for the model first to learn how to start a spreadsheet program from external instruction; then to use that new knowledge as a basis for analogy to discover how to start a graph-drawing package; and then to build on its knowledge by learning through exploration how to draw a graph. The model works with a limited working memory, employs recognition-based problem solving (Howes, 1993), and acquires display-based skill (Payne, 1991) in an interactive, situated task.

These problems of reusability are even more acute when creating models for synthetic environments because of the size and type of models. This is true for several reasons: the knowledge is more extensive and exact than many laboratory domains previously studied. The models must interact with complex, interactive simulations. The work may be classified, which will add an additional constraint in hiring someone with multiple skills. Scenarios may simulate hours of behavior rather than the minutes of typically modeled laboratory tasks. This represents a lot of knowledge, and the timeframe can make troubleshooting more difficult. Finally, there are many cases where an explanation facility is required to explain the model's behavior for other observers.

3.5 Summary

A framework to assist with integration and reuse will have to be developed. It should be common in the sense that the appropriate simulation entities and analysis tools would be available, and for a given application or analysis, developers would plug them together. The DIS protocol and ModSAF are being used in this way to some extent, but they are hard to use and do not support the desired level of ease of use nor the level of cognitive realism.

CHAPTER 4

Current Objective: Improved Usability

In addition to improving the match of synthetic forces to human behavior itself, there are several aspects of these models that must be improved so they can be developed, tested, and used by modelers and analysts. A large amount of time is often required to build models and understand their behavior, more than we believe should be necessary. The difficulties of simply creating and manipulating models of behavior can preclude us from spending more time developing and testing models, and using these models in training or for performing “what-if” analyses.

While Pew and Mavor (1998, p. 10) initially note that their report will not address usability, they later (p. 282) note the need to have quickly reconfigurable models. They also discuss (p. 292) ease of use. This revision is completely appropriate because usability is important. Models that are too difficult to be used are not used. This issue is also being taken up in the next generation of simulation models in the United States (Ceranowicz, 1998).

4.1 Usability of the Models

As we have noted before (Ritter, Jones, & Baxter, 1998b; Ritter & Larkin, 1994), cognitive models suffer from usability problems. Few lessons from the field of Human-Computer Interaction (HCI) have been re-applied to increase the understanding of the models themselves, even though many results and techniques in HCI have been discovered using cognitive modeling.

Modelers have to interact with the model several times and in several ways over the lifetime of the model. As a first step, the models must be easy to create. As part of the creation and validation process, the models must be debugged on the syntactic level (will it run?), on the knowledge level (does it perform the task?), and on a behavioral level (does it perform the task like a human?). All of these levels are important if the costs of acquiring behaviors are to be reduced. While we can point to some recent advances in usability (Anderson & Lebiere, 1998; Jones, 1999b; Kalus & Hirst, 1999; Ritter et al., 1998b), further work will be required.

It is also probably fair to say that cognitive models can often be difficult to explain and understand. This problem has been noted as a result in a recent Air Force model comparison exercise, AMBR, covered in more detail in Section 6.2.7 (Gluck & Pew, 2001a). The difficulty in understanding a model’s behavior is partially due to their complexity, but it is compounded at times by the difficulty of their interfaces not supporting the models in a structured way, not displaying the model’s state, and not supporting exploration of the model’s state. In many cases this is not intentional, but arises out of the modeling languages youth as programming languages, and that support for usability takes time away from applications and modeling itself.

4.2 Desired Accuracy of the Models

Another problem is knowing when to stop improving the model. In science for science's sake, there is no limit—the model is continually improved. In the case of engineering-like applications, such as behavioral models in synthetic environments, knowing when to stop is a valid question. In many cases we do not know how accurate these models have to be in order to be useful and at what point additional accuracy is no longer worthwhile. For example, does having an emotional, simulated opponent lead to better or worse training?

The purpose and goals of each modeling project will help determine when to stop development, so they need to be carefully laid out when developing a model of behavior. The stopping rule also applies to the synthetic environment as well as the model—there is no point in developing a simulation that is too detailed. This question is becoming more important as the models become more accurate and modifiable.

4.3 Aggregation and Disaggregation of Behaviors

A clear requirement for simulations in synthetic environments is the ability to aggregate or summarize subunits and, in other situations, the ability to disaggregate and place the subunits from a larger grouping. When the tanks in a platoon are each simulated in a platform-level simulation, they must be aggregated to display them as a platoon on a more abstract or larger-scale map. Similarly, higher-level units may have to be placed into a simulation when moving a larger unit into a platform-level simulation. This aggregation (or disaggregation) may need to occur multiple times when crossing levels of resolution to provide the right level for a report.

This area has received a limited amount of study, yet it is a common need across multiple types of simulations. None of the cognitive architectures examined in Pew and Mavor (1998, Table 3.1) or here offer any insight. We can only note that several of the architectures (e.g., COGNET, Soar) are designed to support multiple agents.

4.4 Summary

Environments for interacting with existing modeling architectures are generally poorer than those now provided for most programming languages. The requirements for modeling are greater than general programming, including the need for adjustable accuracy, different levels of analyses, and multiple measurements from running programs. These factors contribute to making modeling difficult. We need new models and new techniques for building and using models.

CHAPTER 5

Recent Developments for Modeling

In addition to the architectures and approaches identified by Pew and Mavor (1998), there are a few other architectures that are worth examining. In this chapter we note them, including the lessons they provide. Our reviews also explicitly consider ease of use (i.e., model populating).

We focus our comments on cognitive architectures because they have been created for modeling the strengths and limitations of human behavior. Any system built for other reasons that was adapted in this way—for example, other AI systems—would start to approach these systems in capabilities and limitations. It is quite likely that the cognitive architecture that best matches human behavior will vary by the type of behavior and level of aggregation. For example, different architectures will be preferred for modeling a soldier performing simple physical tasks than for a deliberate and reflective commander.

There will continue to be a range of architectures created. We agree completely with Pew and Mavor (pp. 110-111) that further work is necessary before settling on an architecture. That is not to say that architectures will not continue to converge (e.g., Soar and EPIC, Chong, 2001, and Soar and ACT-R, Jones, 1998). We start, however, by examining ways to summarize data and some advanced AI techniques to help create models. We then examine several architectures.

5.1 Data Gathering and Analysis Techniques

Scattered throughout Pew and Mavor (e.g., pp. 323-325) are comments about the need for data to develop and test models. Data to develop models can come from a wide variety of sources. Data can come from speaking to experts and having them do tasks off-line, so-called knowledge acquisition (Chipman & Meyrowitz, 1993; Schraagen, Chipman, & Shalin, 2000; Shadbolt & Burton, 1995). Data can also come from having experts talk aloud while performing the task (Ericsson & Simon, 1993). Talking aloud is a more accurate way to acquire the knowledge because it is based on actual behavior rather than someone's impression and memory of behavior. It is, however, a more costly approach because the modeler must infer the behavior generators. Data for developing models can also come from non-verbal measurements of experts while they perform the task. Non-verbal measurements are probably the least useful data (but still useful in some circumstances) for developing models. These data are useful, however, in testing models that make timing predictions. Data can also come from previously run studies, reviews, and compendia of such studies (e.g., Boff & Lincoln, 1988; SeKular & Blake, 1994). A useful review of data types and analysis methods in this area is provided by Hoffman (1987).

A major requirement will be a balance between the experimental control of the lab and the richness of the real world. An appropriate balance can sometimes be achieved by gathering data in the same micro-world simulations in which the models will be deployed,

such as synthetic environments. These environments can be used to model all the salient aspects of the real world, while still providing some level of experimental control.

Once the data are in hand, they will often have to be aggregated or summarized. Expert summaries from knowledge acquisition already represent summarized data, but the field of verbal protocol analysis has developed a wide range of techniques for summarizing such data.

Reviews and suggestions in this area are available (Ericsson & Simon, 1993; Sanderson & Fisher, 1994), but there exists a very wide range of techniques that vary based on how advanced the theory is, the purposes of the research, and the domain. Survival analysis is one example of an advanced technique to examine protocol data for temporal patterns for later inclusion and comparison against model behavior (Kuk, Arnold, & Ritter, 1999).

With data in hand, the next step is either to develop a model or to test an existing model. There is little formal methodology about how to create models. Some textbooks attempt to teach this creative task either directly (vanSomeren, Barnard, & Sandberg, 1994) or by example (McClelland & Rumelhart, 1988; Newell & Simon, 1972). There are summaries of the testing process (Ritter & Larkin, 1994) and of some possible tests (Ritter, 1993a). Tenney and Spector (2001); and Ritter and Bibby (2001) provide particularly useful example sets of comparisons. Repairing a model based on the results of the tests can be a task requiring a lot of creativity.

5.2 Advanced AI Approaches

There are some existing AI tools that could be used to create, augment, or optimize models of performance. We note here three tools with which we are particularly familiar. These include approaches for creating behaviors, such as genetic algorithms and traditional AI-planning programs.

5.2.1 Genetic Algorithms

Genetic Algorithms (GAs) are search methods that can be used in domains in which no heuristic knowledge is available and an objective function exhibits high levels of incoherence (Goldberg, 1989). That is to say, a small change to the solution state may often result in large changes to the objective function or fitness measure. These algorithms are expensive in machine resources and exhibit slow (but often steady) convergence to a solution. They might be used as a search strategy of last resort for plan formation.

GAs are a family of algorithms loosely based on Darwinian evolution. They optimize functions without assuming that the search space will be linear. They start with a population of templates for possible solutions (analogous to sets of chromosomes), and evaluate them to determine how well they perform (fitness). After the fitness values are computed, a new population is created. A variety of methods have been used to create the next generation, but in each case the underlying principle has been to include copies of the chromosomes proportional to their fitness, and at each generation to create new combinations by combining two parents' chromosomes. The cycles of evaluation and creation are then repeated.

Heuristics can be used with GAs to seed the initial population in a non-random way or to guide the crossover process in a way that changes the distribution of offspring. Using heuristics results in a *memetic* algorithm (one that manipulates basic blocks of information or *memes*). As has been common experience throughout the history of AI, this introduction of domain knowledge can drastically transform the performance of the GA. Such algorithms have been found to exceed the performance of previous approaches in a number of domains (Burke, Elliman, & Weare, 1995). There may be reason for using GAs as a search strategy in planning.

5.2.2 Tabu Search

Tabu search, as developed by Glover (Glover & Laguna, 1998), is a general purpose approach remarkably effective for difficult problems where the objective function has some local coherence. It is surprising how often hill-climbing approaches such as the A* algorithm are used in current plan-building algorithms, despite the domains being prone to local maxima. Tabu search uses the novel concept of *recency* memory to prevent moves in a solution space from being tried when some component of that state has recently been changed in a previous move. This surprisingly simple idea forces the search away from a local maximum. Long-term memory is used to hold the best solution state found so far and this knowledge may be used to restart the search far away from any previous exploration of the state space.

The Tabu search approach would almost certainly lead to improved solutions with reasonable computational complexity. It would be worth using this approach to search for strategies and plans at various levels in a synthetic environment from the individual combatant to the highest level source of command and control.

Soar is impressive in its ability to reuse parts of problems that have been solved in the past and to plan in a goal-directed way that can seem ingenious. Real human problem solving can be less structured, however, and can leap from one approach to another in a manner that is difficult to model. Tabu search has this characteristic, however, as part of its diversification strategy. Including Tabu search in a cognitive architecture would be interesting. There may be some advantages to be gained by grafting on other similar systems that modify the beliefs of a cognitive architecture so as to maintain various types of logical consistency in the set of facts held.

5.2.3 Multiple Criteria Heuristic Search¹

Heuristic search, one of the classic techniques in AI, has been applied to a wide range of problem-solving tasks including puzzles, two-player games, and path-finding problems. A key assumption of all problem-solving approaches based on utility theory, including heuristic search, is that we can assign a single utility or cost to each state. This, in turn, requires that all criteria of interest can be reduced to a common ratio scale.

¹This section was drafted by Brian Logan and revised by the authors.

The route-planning problem has conventionally been formulated as one of finding a minimum-cost (or low-cost) route between two locations in a digitized map, where the cost of a route is an indication of its quality (e.g., Campbell, Hull, Root, & Jackson, 1995). In this approach, planning is regarded as a search problem in a space of partial plans, allowing many of the classic search algorithms such as A* (Hart, Nilsson, & Raphael, 1968) or variants such as A*epsilon (Pearl, 1982) to be applied. However, while such planners are complete and optimal (or optimal to some bound ϵ), formulating the route-planning task in terms of minimizing a single criterion is difficult.

For example, consider the problem of planning a route in a complex terrain of hills, valleys, impassable areas, and so on. A number of factors will be important in evaluating the quality of a plan: the length of the route, the maximum negotiable gradient, the degree of visibility, and so on. In any particular problem, some of these criteria will affect the feasibility of the route, while others are simply preferences. Route planning is an example of a wide class of multi-criteria, problem-solving tasks, where different criteria must be traded off to obtain an acceptable solution.

One way of incorporating multiple criteria into the problem-solving process is to define a cost function for each criterion and use, for example, a weighted sum of these functions as the function to be minimized. We can, for example, define a *visibility cost* for being exposed and combine this cost with cost functions for the time and energy required to execute the plan to form a composite function that can be used to evaluate alternative plans. However, the relationship between the weights and the solutions produced is complex in reality, and it is often unclear how the different cost functions should be combined linearly as a weighted sum to give the desired behavior across all magnitude ranges for the costs. This makes it hard to specify what kinds of solutions a problem-solver should produce and hard to predict what a problem solver will do in any given situation; small changes in the weight of one criterion can result in large changes in the resulting solutions. Changing the cost function on a single criterion to improve the behavior related to that criterion often leads to changing all the weights for all the other costs as well because the costs are not independent. Moreover, if different criteria are more or less important in different situations, we need to find sets of weights for each situation.

The desirability of trade-offs between criteria is context-dependent. In general, the properties that determine the quality of a solution are incommensurable. For example, the criteria may only be ordered (on an ordinal scale), with those criteria that determine the feasibility of a solution being greatly preferred to those properties that are merely desirable. It is difficult to see how to convert such problems into a multi-criterion optimization problem without making ad hoc assumptions. It is also far from clear that human behavior solely optimizes on a single criterion.

Rather than attempt to design a weighted-sum cost function, it is often more natural to formulate such problems in terms of a set of constraints that a solution should satisfy. We allow constraints to be prioritized, that is, it is more important to satisfy some constraints than others, and soft, that is, constraints are not absolute and can be satisfied to a greater or lesser degree. Such a framework is more general in admitting both optimization problems (e.g., minimization constraints) and satisficing problems (e.g., upper-bound constraints), which cannot be modeled by simply minimizing weighted-sum cost functions. Vicente

(1998) suggests ways in which such constraints can be analyzed as part of a work domain analysis.

This approach to working with constraints provides a way for more clearly specifying problem-solving tasks and more precisely evaluating the resulting solutions. There is a straightforward correspondence between the *real problem* and the constraints passed to the problem-solver. A solution can be characterized as satisfying some constraints (to a greater or lesser degree) and only partially satisfying or not satisfying others. By annotating solutions with the constraints they satisfy, the implications of adopting or executing the current best solution are immediately apparent. The annotations also facilitate the integration of the problem-solver into the architecture of an agent or a decision-support system (see for example, Logan & Sloman, 1998). If a satisfactory solution cannot be found, the degree to which the various constraints are satisfied or violated by the best solution found so far can be used to decide whether to change the order of the constraints, relax one or more constraints, or even redefine the goal, before making another attempt to solve the problem.

The ordering of constraints blurs the conventional distinction between absolute constraints and preference constraints. All constraints are preferences that the problem-solver will try to satisfy, trading off slack on a more important constraint to satisfy another, less important constraint.

The A* search algorithm is ill-suited to dealing with problems formulated in terms of constraints. Researchers at Birmingham have therefore developed a generalization of A* called A* with Bounded Costs (ABC; Alechina & Logan, 1998; Logan & Alechina, 1998), which searches for a solution that best satisfies a set of prioritized soft constraints.

The utility of this approach and the feasibility of the ABC algorithm have been illustrated by an implemented route planner that is capable of planning routes in complex terrain satisfying a variety of constraints. This work was originally motivated by difficulties in applying classical search techniques to agent-route planning problems. However, the problems identified with utility-based approaches, and the proposed solutions, are equally applicable to other search problems.

5.3 Psychologically Inspired Architectures

We review here several psychologically inspired cognitive architectures that were not covered by Pew and Mavor (1998). These architectures are interesting because (1) they are psychologically plausible, (2) some of them provide examples of how emotions and behavioral moderators can be included, and (3) several illustrate that better interfaces for creating cognitive models are possible.

5.3.1 Elementary Perceiver and Memoriser

The Elementary Perceiver And Memoriser (EPAM) is a well-known computer model of a wide and growing range of memory tasks. The basic ideas behind EPAM include mechanisms for encoding chunks of information into long-term memory by constructing a discrimination network. The EPAM model has been used to simulate a variety of psychological regularities, including the learning of verbal material (Feigenbaum & Simon,

1962, 1984) and expert digit-span memory (Richman, Staszewski, & Simon, 1995). EPAM has been expanded to use visuo-spatial information (Simon & Gilmarin, 1973).

EPAM organizes memory into a collection of chunks, where each chunk is a meaningful group of basic elements. For example, in chess, the basic elements are the pieces and their locations; the chunks are collections of pieces, such as a king-side pawn formation. These chunks are developed through the processes of discrimination and familiarization. Essentially, each node of the network holds a chunk of information about an object in the world. The nodes are interconnected by links into a network with each link representing the result of applying a test to the object. When trying to recognize an object, the tests are applied beginning from the root node, and the links are followed until no further test can be applied. When a node is reached, if the stored chunk matches that of the object then familiarization occurs. The chunk's resolution is then increased by adding more details of the features in that object. If the current object and the chunk at the node reached differ in some feature, then discrimination occurs, which adds a new node and a new link based on the mismatched feature. Therefore, with discrimination, new nodes are added to the discrimination network; with familiarization, the resolution of chunks at those nodes is increased.

The Chunk Hierarachy and REtrieval STructures (CHREST; de Groot & Gobet, 1996; Gobet & Simon, 1996b) is one of the most current theories of memory developed from the ideas in EPAM. Gobet and Simon (2000) present a detailed description of the present version of CHREST and report simulations on the role of presentation time in the recall of game and random chess positions. As in the earlier chunking theory of Chase and Simon (1973), CHREST assumes that chess experts develop a large EPAM-like net of chunks during their practice and study of the game. In addition, CHREST assumes that some chunks, which recur often during learning, develop into more complex retrieval structures (templates) with slots for variables that allow a rapid encoding of chunks or pieces.

EPAM and its implementations are important to consider because they fit a subset of regularities in memory very well. This at least serves as an example for other theories and architectures to emulate. It may also be possible to include the essentials of EPAM in another system, such as Soar or ACT-R, extending the scope of both approaches.

5.3.2 Neural Networks

Pew and Mavor (1998, chap. 3) review neural networks. Here, therefore, we only provide some further commentary, introduce some more advanced concepts, and note a few further applications.

Connectionist systems have demonstrated the ability to learn arbitrary mappings. Architectures such as the Multi-Layer Perceptron (MLP) are capable of being used as a black box that can learn to recognize a pattern of inputs as a particular situation. This requires supervised training and may involve heavy computational resources to arrive at a successful solution using the back-propagation algorithm. Training can be continued during performance as a background task, and thus, an entity could have an ability to learn during action based on this approach. Recognition performance is relatively rapid and a multi-layer perceptron might be used to model a reaction mechanism in which a combatant responds to

coming under fire, or spotting the presence of the enemy, for example. It might also be used to activate particular aspects of military doctrine depending on the current circumstances.

Recurrent nets such as the Elman (1991) net have the ability to generate sequences of tokens as output. These seem to offer some promise of detecting an input situation and producing a series of behavioral actions as a response. This behavior of recurrent nets may be useful for modeling the reactive behavior of an entity over a short time period, while a symbolic cognitive model is used for the higher-level cognitive processes that occur over a longer time span.

5.3.3 Sparse Distributed Memories

Subtle issues such as the *tip-of-the-tongue* phenomena (Koriat & Lieblich, 1974) and the fact that we know if we know something (feeling of knowing) before becoming aware of the answer are not often modeled (although, see Schunn, Reder, Nhouyvanisvong, Richards, & Stroffolino, 1997, for a counter example). These effects may be captured using memory models such as Kanerva's (1988) Sparse Distributed Memory (SDM), which has been put forward as a plausible model of brain architecture, particularly the cerebellum, as well as by Albus's (1971) Cerebellar Model Arithmetic Computer (CMAC).

The way in which a combatant's experience of the world is stored and modeled is important. An SDM seems to offer powerful human-like ways of recalling nearest matches to present experience in a best-first manner. They have the interesting property of storing memories such that recall works by finding the best match to imperfect data. They are also a natural way of storing sequences. They exploit interesting mathematical properties of binary metric spaces with a large number of dimensions. It is intriguing that SDMs have the human-like properties that they "know if they know" something before the retrieval process is complete. They also exhibit the tip-of-the-tongue phenomenon and replicate the human ability to recall a sequence or tune given the first few items or notes. They can also learn actuator sequences that might be used in muscle control or reflex patterns of behavior. This can even be seen as a kind of thinking by analogy that has a uniquely human-like ability to find a close match rapidly without exhaustive or even significant time spent in search.

5.3.4 PSI and Architectures That Include Emotions

PSI is a relatively new cognitive architecture designed to integrate cognitive processes, emotions, and motivation (Bartl & Dörner, 1998). The architecture includes six motives (needs for energy, water, pain avoidance, affiliation, certainty, and competence). Cognition is modulated by these motive/emotional states and their processes. In general, PSI organizes its activities similar to Rasmussen's (1983) hierarchy: first, it tries highly automatic skills if possible, then it skips to knowledge-based behavior, and as its *ultima ratio* approach it uses trial-and-error procedures. It is one of the only cognitive architectures that we know about that takes modeling emotion and motivation as one of its core tasks. Its source code, in Delphi Pascal, is available (www.uni-bamberg.de/ppp/insttheopsy/psi-software.html).

A model in the PSI architecture has been tested against a set of data taken from a dynamic control task. The model performed the same task and its number of control actions was within the range of human behavior. Its predictions of summary scores were outside the

range of human behavior—the model was less competent (Detje, 2000)—but single subjects can be modeled by varying starting parameters (Dörner, 2000). In such a complex task as the “Island” scenario some people will use meta-cognition to improve their performance (particularly if they are encouraged to think aloud as they were in Detje’s study). The same data could reveal that these subjects profit from meta-cognition and that their behavior then differs from what is implemented currently in PSI (see Bartl, 2000, for a more detailed explanation).

This model needs to be improved before it matches human emotional data as well as other cognitive models match non-emotional data. It is, however, one of the few models of emotion compared with data.

The PSI architecture is currently incomplete, which raises interesting questions about how to judge a nascent architecture. PSI does not have a large enough user community and has not been developed long enough to have a body of regularities to be compared with let alone adjusted to fit. How can PSI be compared with the older architectures with existing tutorials, user manuals, libraries of models, and example applications?

Several other models of emotions and architectures that use emotions have been created. Reviews of emotional models (Hudlicka & Fellous, 1996; Picard, 1997) typically present models and architectures that have not been compared and validated against human data. There appears to be one other exception, an unpublished PhD thesis by Araujo at the University of Sussex (cited in Picard, 1997). Some of us are attempting to add several simple emotions to ACT-R (Belavkin, 2001; Belavkin et al., 1999) and validate the model by comparing the revised model with an existing model and comparable data (G. Jones, Ritter, & Wood, 2000).

5.3.5 COGENT

COGENT is a design environment for creating cognitive models and architectures (Cooper & Fox, 1998). It allows the user to draw box-and-arrow diagrams to structure and illustrate the high-level organization of the model and to fill in the details of each box using one or a series of dialogue sheets. The boxes include inputs, outputs, memory buffers, processing steps, and even production systems as components.

COGENT’s strengths are that it is easy to teach, the displays provide useful summaries of the model that help with explanation and development, and the environment is fairly complete. It appears possible to reuse components on the level of boxes. COGENT’s weaknesses are that it is fairly unconstrained; for large systems it may be unwieldy; and it might not interface well to external simulations.

COGENT also shows that cognitive modeling environments can at least appear more friendly. The results of its graphic interface routinely appear in talks as model summaries. The interface is also quite encouraging to users, allowing them to feel that they can start working immediately.

5.3.6 Hybrid Architectures

Hybrid architectures are architectures that typically include symbolic and non-symbolic elements. A more general definition would be architectures that include major components from multiple architectures.

Hybrid architectures are mentioned briefly by Pew and Mavor (1998, pp. 108-110). Work has continued in this area with some interesting results. LICAI (Kitajima & Polson, 1996; Kitajima, Soto, & Polson, 1998), for example, models how people explore and use interfaces based on a theory of how Kintsch's (1998) schemas receive activation. The U.S. Office of Naval Research (ONR) has sponsored a research program on hybrid architectures (Gigley & Chipman, 1999). This has given rise to some new, interesting hybrid architectures (e.g., Sun, Merrill, & Peterson, 1998; Wang, Johnson, & Zhang, 1998).

Perhaps the most promising hybrids are melding perception components across cognitive architectures. The EPIC (Kieras & Meyer, 1997) architecture's perception and action component has been merged with ACT-R's perceptual-motor component, ACT-R/PM (Byrne, 2001; Byrne & Anderson, 1998) and with Soar (Chong, 2001). This has led to direct reuse and unification. Similar results have been found with the Nottingham functional interaction architecture being used by Soar and ACT-R models (Bass et al., 1995; Baxter & Ritter, 1996; Ritter et al., 2000; G. Jones et al., 2000).

5.4 Knowledge-Based Systems and Agent Architectures

Agent architectures will be important within synthetic environments for modeling autonomous vehicles and for exploring the doctrine of autonomous vehicles. Most principled agent architectures have historical roots in distributed artificial intelligence. For several decades, distributed AI has been tackling essentially the same problem as Knowledge-Based Systems (KBS) research, namely, how to produce efficient problem-solving behavior in software. The main concept that brings agency and KBS together is the idea of operation at the knowledge level as described by Newell (1982).

The behavioral law used by an observer to understand the agent at the knowledge level is the principle of maximum rationality (Newell, 1982), which states, "If an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action." The modeling of intelligent artificial systems at the knowledge level, that is, with no reference to details of implementation, is a key principle in KBS construction. It is also at the heart of many assumptions in the tradition of explaining human behavior.

Nwana (1996) claims that an important difference between agent-based applications and other distributed computing applications is that agent-based applications operate typically at the knowledge level, whereas distributed computing applications operate at the symbol level. At the symbol level, the entity is seen simply as a mechanism acting over symbols, and its behavior is described in these terms.

The theoretical links between the motivations behind KBS and agent research can be seen in the main approaches taken to the definition of software agency. Ascriptional agency attempts to create convincing human-like behaviors in software in the belief that this will produce programs that are easy to interact with. This work can be seen as paralleling the expert behavioral modeling approach that is currently widely espoused in the KBS

community. The Belief-Desire-Intention (BDI) agents focus on the concept of intentionality—the mental attitudes of the agent. BDI models have been successfully implemented in systems such as the DESIRE framework (Brazier, Dunin-Keplicz, Treur, & Verbrugge, 1999) and the JAVA Agent Compiler and Kernel (JACK) component system (Busetta, Howden, Rönnquist, & Hodgson, 1999a; Busetta, Rönnquist, Hodgson, & Lucas, 1999b).

JACK is an extension to JAVA. It includes a JAVA library and a compiler that takes a JAVA program with embedded JACK statements. A JAVA compiler expands/incorporates the JACK statements to create a runnable JAVA program. These statements implement a BDI architecture, while allowing JAVA statements to extend and implement them. The statements include commands like `@achieve(condition, event)`, which subgoals on *event* if *condition* is not found to be true.

The resulting program instantiates a BDI agent. Its BDI architecture is made up of beliefs represented with a database; desires represented as events that can trigger plans; and intentions represented through these plans. For example, a fact may come in from perception and match a desire, that of putting new facts into the database. This may result in further desires being matched and intentions (plans) leading to behaviors. Further information is available at the JACK developer's website (www.agent-software.com.au).

Reviews of the agent literature (Etzioni & Weld, 1995; Franklin & Graesser, 1997; Wooldridge & Jennings, 1995)² reveal that, when attempting to define agency as dependent on the possession of a set of cardinal attributes, many of the attributes suggested could also be seen as characteristic of behavior that is best explained at the knowledge level. These include abstraction and delegation, flexibility and opportunism, task orientation, adaptivity, reactivity, autonomy, goal-directedness, flexibility, collaborative and self-starting behavior, temporal continuity, knowledge-level communication ability, social ability, and cooperation.

Both agent systems and KBSs are moving in the direction of modular components of expertise as a response to the problems of knowledge use and reuse to promote intelligent behavior in software. Domain ontologies form a significant subset of these KBS components. Increasingly, multi-agent systems are being produced that use such domain ontologies to facilitate agent communication at the knowledge level, for example, the agent network created as part of the Infosleuth architecture (Jacobs & Shea, 1996). Some agent systems also draw explicitly on models of problem-solving expert behavior developed in KBS research. The internet-based Multi-agent Problem Solving (IMPS) architecture (Crow & Shadbolt, 1998) uses software agency as a medium for applying model-driven knowledge engineering techniques to the internet. It involves software agents that can conduct structured online knowledge acquisition using distributed knowledge sources. Agent-generated domain ontologies are used to guide a flexible system of autonomous agents driven by problem-solving models.

² For online information about examples and related U.S. programs, see www.darpa.mil/ito/ResearchAreas.html and www.nosc.mil/robots/air/amgsss/mssmp.html.

5.5 Architectural Ideas Behind the Sim_Agent Toolkit³

Since the early 1970s, Sloman and his colleagues have been attempting to develop requirements and designs for an architecture capable of explaining a wide variety of facts about human beings and other intelligent agents. Sloman's ideas about cognitive architectures and the theoretically based agent architecture toolkit (Sim_Agent) provide useful lessons about architectural toolkits and about process models of emotions. Further information is available at the CogAff website (www.cs.bham.ac.uk/~axs/cogaff.html).

5.5.1 Cognition and Affect

A human-like information processing architecture includes many components performing different functions all of which operate in parallel, asynchronously. This is not the kind of low-level parallelism found in neural nets (although such neural mechanisms are part of the infrastructure). Rather there seem to be many functionally distinct modules performing different sorts of tasks concurrently, a significant proportion of them are concerned with the monitoring and control of bodily mechanisms, for example, posture, saccades, grasping, temperature control, daily rhythms, and so on.

The very oldest mechanisms in the human architecture are probably all reactive in the sense described in various recent papers (e.g., Sloman, 2000). The key feature of reactivity is the lack of "what-if" reasoning capabilities, with all that entails, including the lack of temporary workspaces for representations of hypothesized futures (or past episodes); the lack of mechanisms for stored factual knowledge (generalizations and facts about individuals) to support the generation of possible futures, possible actions, and likely consequences of possible actions; and the lack of mechanisms for manipulating explicit representations.

Both reactive and deliberative mechanisms require perceptual input and can generate motor signals. However, to function effectively, both perceptual and action subsystems may have evolved new layers of abstraction to support the newer deliberative processes, for example, by categorizing both observed objects and events at a higher level of abstraction, and allowing higher-level action instructions to generate behavior in a hierarchically organized manner. More generally, different subsystems use information for different purposes so that a number of different processes of analysis and interpretation of sensory input occur in parallel, extracting different affordances from raw data from the optic array. Recent work by brain scientists on ventral and dorsal visual pathways are but one manifestation of this phenomenon.

The interactions between reactive and deliberative layers are complex and subtle, especially as neither is in charge of the other, though at times either can dominate. Moreover, the division is not absolute: information in the deliberative system can sometimes be transferred to the reactive system (e.g., via drill and practice learning), and information in the reactive system can sometimes be decompiled and made available to deliberative mechanisms (though this is often highly error-prone).

³ This section was drafted by Aaron Sloman and revised by the authors.

For reasons explained in various papers available in the CogAff FTP site, it is possible to conjecture that at a much later evolutionary stage a third class of mechanism developed, again using and redeploying mechanisms that had existed previously. The new type of mechanism, which has been provisionally labeled “meta-management,” provides the ability to do for internal processes what the previous mechanisms did for external processes: namely it supports monitoring, evaluation, and control of other internal processes, including, for instance, thinking about how to plan, or planning better ways of thinking. For example, a deliberative system partly driven by an independent reactive system and sensory mechanisms can unexpectedly acquire inconsistent goals. A system with meta-management can notice and categorize such a situation, evaluate it, and perhaps through deliberation or observation over an extended period, develop a strategy for dealing with such conflicts.

Similarly, meta-management can be used to detect features of thinking strategies and, perhaps in some cases, notice flaws or opportunities for improvement. Such a mechanism (especially in conjunction with an external language) also provides a route for absorption of new internal processes from a culture, thereby allowing transmission between generations of newly acquired information without having to wait for new genetic encodings of that information to evolve. Through internal monitoring of sensory buffers, the extra layer adds a kind of self-awareness that has been the focus of discussions of consciousness, subjective experience, qualia, etc. As with external processes, the monitoring, evaluation, and re-direction of internal processes is neither perfect nor total and, as a result, mistakes can be made about what is going on, inappropriate evaluations of internal states can occur, and attempts to control processing may fail, for example, when there are lapses of attention despite firm intentions.

Another feature of meta-management is its ability to be driven by different collections of beliefs, attitudes, strategies, and preferences, in different contexts, explaining how a personality may look different at home, driving a car, in the office, etc. Besides the three main concurrent processing layers (reactive, deliberative, and meta-management) identified above that others have found evidence for, a number of additional specialized mechanisms are needed, including: mechanisms for managing short- and long-term goals, a variety of long- and short-term memory stores, and one or more global alarm systems capable of detecting a need for rapid global re-organization of activity (freezing, fleeing, attacking, becoming highly attentive, etc.), and also producing that re-organization.

For instance, whereas many people have distinguished primary and secondary emotions (e.g., Damasio, 1994), Sloman and his colleagues have proposed a third type, tertiary emotions, sometimes referred to as *perturbances* (Sloman, 1998a; Sloman & Logan, 1999). Primary emotions rely only on the reactive levels in the architecture. Secondary emotions require deliberative mechanisms. Tertiary emotions are grounded in the activities of meta-management, including unsuccessful meta-management. There are other affective states concerned with global control, such as moods, which also have different relationships to the different layers of processing. Many specific states that are often discussed but very unclearly defined, such as arousal, can be given much clearer definitions within the framework of an architecture that supports them.

It looks as if various subsets of the capabilities described here arising out of the three layers and their interactions can be modeled in the architectures developed so far, for example, Soar, ACT-R/PM, Moffatt and Frijda’s Will architecture (2000), and the various

logic-based models that dominate the ATAL (Architectures, Theories and Languages) series of workshops (e.g., Wooldridge, Mueller, & Tambe, 1996, also see mas.cs.umass.edu/atal/), and books like Wooldridge and Rao (1999).

However, only small subsets of these capabilities can be modeled at present. Any realistic model of human processing needs to be able to cope with contexts including rich bombardment with multi-modal sensory and linguistic information; where complex goals and standards of evaluation are constantly interacting; where things often happen too fast for fully rational deliberation to be used; where everything that occurs does not always fall into a previously learned category for which a standard appropriate response is already known; where decisions have to be taken on the basis of incomplete or uncertain information; and where the activity of solving one problem or carrying out one intricate task can be subverted by the arrival of new factual information, new orders, or new goals generated internally as a side-effect of other processes.

Where the individual is also driving a fast-moving vehicle or is under fire then it is very likely that a huge amount of the processing going on will involve the older reactive mechanisms, including many concerned with bodily control and visual attention. It may be some time before we fully understand the implications of such total physical immersion in stressful situations, including the effects on deliberative and meta-management processes. (For example, fixing attention on a hard planning problem can be difficult if bombs are exploding all around you. Can our models explain why?)

5.5.2 Sim_Agent and CogAff

Sloman and his colleagues' general architectural toolkit, the Sim_Agent Toolkit, allows them to explore a variety of new ideas about complex architectures. It is not an architecture, but a steadily developing toolkit for exploring architectures.

The CogAff architecture provides a schema, based on a 3 by 3 grid that provides a framework for describing specific architectures according to the grid components present, their control relationships, and how information flows between them. H-CogAff is a specific human-like version that is a particularly rich special case. Other special cases include various kinds of purely reactive (i.e., non-deliberative) architectures (perhaps insects or reptiles), Brooks' subsumption architectures, purely deliberative architectures (lots of old AI systems, early versions of Soar and ACT), and so on.

Sloman and his colleagues also wanted a toolkit that supported exploration of a number of interacting agents (and physical objects, etc.) where each agent has a variety of very different mechanisms running concurrently and asynchronously, yet influencing one another. They also wanted to be able to very easily change the architecture within an agent, change the degree and kind of interaction between components of an agent, and speed up or slow down the processing of one or more sub-mechanisms relative to others (Sloman, 1998b). In particular, they wanted to be able to easily combine different types of symbolic mechanisms and also sub-symbolic mechanisms within one agent. The toolkit was also required to support rapid prototyping and interactive development with close connections between internal processes and graphic displays.

Because other toolkits did not appear to have the required flexibility and tended to be committed to a particular type of architecture, Sloman and his colleagues built their own

toolkit, which has been used for some time at the University of Birmingham and DERA, Malvern. Their toolkit is described briefly in Sloman and Logan (1999) and in more detail in the online documentation at the Birmingham Poplog FTP site (<ftp.cs.bham.ac.uk/pub/dist/poplog/>). The code and documentation are freely available online. The toolkit runs in Pop-11 in the Poplog system (inherently a multi-language AI system, so that code in Prolog, Lisp, or ML can also be included in the same process). Poplog has become freely available (www.cs.bham.ac.uk/research/poplog/freepoplog.html).

At present, Sloman does not propose a specific overarching architecture as a rival to systems like Soar or ACT-R. He feels that not enough is yet known about how human minds work and, consequently, any theory proposing *the* architecture is premature. Instead, he and his group have been exploring and continually refining a collection of ideas about possibly relevant architectures and mechanisms. Although the ideas have been steadily developing, Sloman and his colleagues do not believe that they are near the end of this process. So, although one could use a label like *CogAff* to refer to the general sort of architecture they are currently talking about, it is not a label for a fixed design. Rather *CogAff* should be taken to refer to a high-level overview of a class of architectures in which many details still remain unclear. The *CogAff* ideas are likely to change in dramatic ways as more is learned about how brains work, about ways in which they can go wrong (e.g., as a result of disease, aging, brain damage, addictions, stress, abuse in childhood, etc.), and how brains differ from one species to another, or one person to another, or even within one person over a lifetime.

The toolkit is still being enhanced. In the short term, they expect to make it easier to explore architectures including meta-management. Later work will include better support for sub-symbolic spreading activation mechanisms and the development of more reusable libraries, preferably in a language-independent form.

5.5.3 Summary

The *Sim_Agent* toolkit and the goals its developers have for it have some commonalities with other approaches. The need for a library of components is acknowledged. They emphasize that reactive behaviors are necessary and desirable, and that the emotional aspects arise out of the reactive mechanisms. It provides a broad range of support for testing and creating architectures. The toolkit provides support for reflection as a type of meta-learning. Other architectures will need to support reflection as well, particularly where the world is too fast-paced for learning to occur during the task (John, Vera, & Newell, 1994; Nielsen & Kirsner, 1994).

The features that the toolkit supports help define a description of architectural types. The capabilities that can be provided, from perception to action and from knowledge to emotion, provide a way of describing architectures.

The major drawback is that none of the models or libraries created in *Sim_Agent* have been compared with human data directly. In defense of this lack of comparison, Sloman claims that the more complex and realistic an architecture becomes, the less sense it makes to test it directly. Instead, he claims that the architecture has to be tested by the depth and variety of the phenomena it can explain, like advanced theories in physics, which also cannot be tested directly.

5.6 Engineering-Based Architectures and Models

There is a history of studying process control in and near industrial engineering that includes studying human operators. This approach is not (yet) part of mainstream psychology, and Pew and Mavor (1998) do not make many references to work in this field.

If tank operators and ship captains can be viewed as running a process, and we believe they can, there is a wide range of behavioral regularities referenced and modeled in engineering psychology that can be generalized and applied to other domains. Major contributions in this area include Reason's (1990) book on errors, Rasmussen's skill hierarchy (1983), the Cognitive Reliability and Error Analysis Method (CREAM) methodology for analyzing human performance (Hollnagel, 1998), and numerous studies characterizing the strengths and weaknesses of human operator behavior (de Keyser & Woods, 1990; Sanderson, McNeese, & Zaff, 1994).

Engineers have also created intelligent architectures. These architectures have almost exclusively been used to create models of users of complex machinery, ranging from nuclear power plants to airplanes. The models are often, but not always, tied to simulations of those domains. Their approach is generally more practical. They are more interested in approximate timing and the overt behavior than in detailed mechanisms. These developers appear to be less interested in the internal mechanisms giving rise to behavior as long as the model is usable and approximately correct.

These models of operators include models of nuclear power plant operators, the Cognitive Simulation Model (COSIMO; Cacciabue et al., 1992), and the Cognitive Environment Simulation (CES; Woods, Roth, & Pople, 1987). AIDE (Amalberti & Deblon, 1992) is a model of fighter pilot behavior; the Step Ladder Model or Skill-based, Rule-based, Knowledge-based model is a generally applicable framework, originally formulated in electronics troubleshooting (e.g., Rasmussen, 1983).

We will also look at a few operator models in more detail.

5.6.1 APEX⁴

APEX (Freed & Remington, 2000; Freed et al., 1998; John et al., 2002) is a set of tools for simulating human performance when interacting with interfaces to perform tasks similar to MIDAS (Laughery & Corker, 1997). The main driver for APEX is the need to model behavior in environments, such as air traffic control and commercial jet flight decks, and to help engineers design usable systems in these domains

Powerful action-selection mechanisms of the sort developed by artificial intelligence researchers are used to cope adaptively with time-pressure and uncertainty, and to coordinate the execution of multiple tasks (i.e., strategic multi-tasking). Usability is taken very seriously (Freed & Remington, 2000). A high-level modeling language is included. Applications of APEX have included time-analysis of skilled behavior, partially-automated

⁴ Comments from Michael Freed were helpful in preparing this section.

human-factors design analysis, and creation of artificial human participants in large-scale simulations.

This general approach has proven successful in allowing APEX to automate the CPM-GOMS HCI analysis method (John & Kieras, 1996) and for reconstructing incidents involving human error in a way that promise eventual error-prediction capabilities. As much as it implements CPM-GOMS, it inherits CPM-GOMS' empirical support. Consistent with the needs of domains in which APEX has been most frequently used, the action-selection architecture emphasizes capabilities having to do with multi-task management, especially regarding concurrency control and strategic task management.

APEX was created by Freed as part of his doctoral dissertation and continues to be developed by researchers at NASA Ames Research Center and elsewhere. They are explicitly concerned about building a community of users to share ideas and models. Further information, including APEX itself, is available through search engines.

APEX is probably best described as an engineering model because it has been designed to serve engineering goals. APEX is interesting because it models the whole operator, from perception to action, and the model often interacts with fairly complete and complex simulations, and can make very detailed predictions easily. It does not yet include learning, and the complex results past CPM-GOMS could be tested more, but the full toolset suggests that interface design tools based on cognitive models are now possible.

5.6.2 Simplified Model of Cognition and Contextual Control Model

The Simplified Model of Cognition (SMoC) (Hollnagel & Cacciabue, 1991) is an extension of Neisser's (1976) perceptual cycle and describes cognition in terms of four essential elements: (1) observation/identification, (2) interpretation, (3) planning/selection, and (4) action/execution. Although these are normally linked in a serial path, other links are possible between the various elements. The small number of cognitive functions in SMoC reflects the general consensus of opinion that has developed since the 1950s on the characteristics of human cognition. The fundamental features of SMoC are the distinction between observation and inference (overt vs. covert behavior), and the cyclical nature of cognition (cf. Neisser, 1976).

SMoC was formulated as part of the System Response Generator (SRG) project (Hollnagel & Cacciabue, 1991). SRG was a software tool developed to study the effect of human cognition (specifically actions and decision making) on the evolution of incidents in complex systems.

The Contextual Control Model (CoCoM; Hollnagel, 1993) is an extension of the SMoC, and addresses the issues of modeling both competence and control. In most models the issue of competence is supported by a set of procedures or routines that can be employed to perform a particular task when a particular set of pre-defined conditions obtains. CoCoM further proposes that there are four overlapping modes of control—influenced by knowledge and skill levels—that also influence behavior:

- Scrambled control: where the selection of the next action is unpredictable. This is the lowest level of control.
- Opportunistic control: where the selection of the next action is based on the current context without reference to the current goal of the task being performed.
- Tactical control: where performance is based on some form of planning.
- Strategic control: where performance takes full account of higher-level goals. This is the highest level of control.

The transition between control modes depends on a number of factors, particularly the amount of subjectively available time and the outcome of the previous action. These two factors are interdependent, however, and also depend on aspects such as the task complexity and the current control mode.

CoCoM has been used in the development of the CREAM (Hollnagel, 1998) within the field of human reliability analysis. The CREAM is a method for analyzing human performance when working with complex systems. It can be used in both the retrospective analysis of accidents and events, and in predicting performance for human reliability assessment. Extending the CREAM is presented below as a useful project.

5.6.3 Summary

These engineering-based architectures suggest that engineering models can provide useful behavior even when the internal mechanisms are not fully tested or perhaps even plausible. These architectures suggest that some of the difficulty in creating the architectures is due to the implicit and explicit knowledge that psychologists bring with them regarding plausibility. We believe psychologists' domain knowledge leads to more accurate models but slower development.

5.7 Summary of Recent Developments for Modeling Behavior

This chapter has reviewed several architectures. These architectures and their applications show that it is becoming increasingly possible to create plausible and useful architectures based on a variety of approaches.

An agreed, formal scheme for classifying architectures would be useful. This ideal system classification would note the sorts of tasks that each architecture performs best, supporting users to choose an architecture for a particular task. The best that we have found is Table 3.1 in Pew and Mavor (1998, pp. 98-105). Our Table 5.1 provides a summary of the architectures presented here in that same format as a supplement to their table. We have included all relevant information of which we are aware for each architecture. In most cases the developers of the architectures have helped complete their entry in this table. Another approach for classifying architectures is available from Logan (1998).

Developments in AI continue to be useful. The general AI methods discussed are not included in this table because they are not broad enough to be considered a cognitive architecture, but they are likely to be useful additions to architectures, either directly or indirectly. For example, genetic algorithms have been included in a proposed architecture

(Holland, Holyoak, Nisbett, & Thagard, 1986), and planning algorithms have been included as adjuncts to Soar (Gratch, 1998). These developments will help extend architectures by providing algorithms for inclusion within architectures, particularly hybrid architectures.

There are several interesting trends to note. One is that the diversity of architectures is not decreasing. New, fundamental ideas on which to base architectures has widened from simply problem solving. For example, EPAM is based on pattern recognition, and PSI and architectures created in the Sim_Agent Toolkit are based on ideas about emotions.

Another interesting trend is that some aspects of the architectures are starting to merge and be reused. The interaction aspects of EPIC have been reused by Soar and by ACT-R. The Nottingham Interaction Architecture is similar in some ways and getting similar reuse (e.g., Jones et al., 2000). These strands are becoming quite similar to each other (Byrne, Chong, Freed, Ritter, & Gray, 1999) and are quite likely to merge in the future.

The importance of model usability is becoming more recognized. COGENT provides an example of how easy a modeling tool should be to pick up and use. Similar developments with Soar and ACT-R are starting to emphasize reusable code, better documentation, and better tutorial materials. Other architectures will have to follow suit to attract users and to train and support their existing users. Newell (1990) wrote about the entry level (the bar) being raised as architectures develop through competition. It is interesting that usability is perhaps the first clear comparison level.

Table 5.1: Comparison of Architectures

Architecture	Original purpose	Submodels
		Sensing and perception
1 EPAM	Model high-level perception, learning, and memory	Visual, auditory perceptual discrimination in real-time (assuming feature-based description of objects)
2 SDM	Simulation of cerebellum as a content-addressable memory	Can be used to recall the nearest stored memory to any encoded perceptual input
3 PSI	Explores interaction of cognition, motivation, and emotion to build an integrated model of human action regulation	Optical perception by "Hypercept" process that scans (simulated) environment for basic features. Raises hypotheses about sensory schemas to which features may belong and tests these hypotheses by subsequent scanning of environment (comparable to saccadic eye-movements). If pattern not recognizable, new schema generated
4 COGENT	Design environment for modeling cognitive processes	Input buffers that can be modified to represent vision and hearing
5 JACK as example of BDI architectures	Constitute an industrial-strength framework for agent applications	JAVA methods + inter-agent messaging
6 Sim_Agent Toolkit	Explores architectures using rapid prototyping	Defined by methods for each agent class.
7 Engineering-based models (e.g., APEX)	Provide models of humans in control loops	Varies, but exists for most models

Table 5.1: Comparison of Architectures (continued)

Submodels			
	Working/ Short-Term memory	Long-Term Memory	Motor Outputs
1	4-7 slot STM; in some versions (e.g., EPAM-IV), more detailed implementation of auditory (Baddeley-like) STM & visual STM	Discrimination net. In recent versions, nodes of discrimination net used to create semantic net and productions	Eye movements, simple drawing behavior
2	Not modeled	Sparse Distributed Memory models related to PDP and neural-net memory models	Motor sequences can be learned. Nearest match memories can be sequences that could be behaviors
3	The <i>head</i> of a protocol memory that permanently makes a log of actions and perceptions	The remnants of logs decay with time. Strings of logs associated with need satisfaction or with pain will be reinforced and have a greater chance to survive and form a part of long-term memory than neutral sequences of events	Basic motor patterns (actions) combined to form complex sensory-motor-programs by learning (i.e., by reinforcement of the successful sensory-motor-patterns in logs)
4	Various types supported	Various types supported	Simple buffer representation of commands
5	Object-oriented structures (JAVA), plus relational modeling support (JACK)	All JAVA support including database interfaces etc. Support for data modeling in JAVA and C++ using JACOB (JACK Object Builder)	JAVA methods
6	List structures	List structures, rules, and arbitrary Pop-11 data structures. Can also use neural nets, if required	Defined by methods for each agent class
7	Usually simple, but extant	Usually simple, but extant	Usually extant, but usually not complex

Table 5.1: Comparison of Architectures (continued)

Architecture	Knowledge Representation	
	Declarative	Procedural
1 EPAM	Chunks, schemas (templates); using nodes in discrimination net	Productions using nodes in discrimination net
2 SDM	A sparse set of memory addresses where data <i>are</i> addresses	Memories naturally form sequences that could be considered procedures
3 PSI	Sensory and sensory-motor patterns consisting of pointer structures forming schemas. A schema includes information about more basal elements and relations of elements in space and time, including language patterns pointing to sensory and sensory-motor patterns (implementation in progress)	Sensory-motor-patterns forming automatisms
4 COGENT	Numbers, strings, lists, tuples, connectionist networks	Production rules, connectionist networks, Prolog
5 JACK as an example of BDI architectures	Object-oriented structures (JAVA), plus relational modeling support	JACK plans and JAVA methods
6 Sim_Agent Toolkit	List structures and arbitrary Pop-11 data structures (e.g., could be constrained to express logical assertions but need not be). Could use neural nets or other mechanisms	Rule sets and arbitrary Pop-11 procedures that can also invoke Prolog or external functions
7 Engineering-based models (e.g., APEX)	Varies, but usually simple	Varies, but usually simple. Many use some form of schemas

Table 5.1: Comparison of Architectures Covered (continued)

Higher-Level Cognitive Functions				
	Learning	Planning	Decision Making	Situation Assessment
1	Chunking, creation of schemas, and production learning online (incremental) and stable against erroneous data	Connections between templates used in planning	Knowledge based	Overt and inferred
2	By incrementing weights across a probability distribution	Does not plan, but can remember plans	Iterative memory recall process	Can learn a set of assessments and generalize these
3	Associative and perceptual learning; operant conditioning: sensory-motor learning, learning goals (situations that allow need satisfaction) and aversions (situations or objects that cause needs)	Built-in hill-climbing procedure: action schemata (i.e., sensory-motor-patterns) are recombined to form new plans. If planning unsuccessful or impossible due to lack of information, trial-and-error procedures used to collect environmental information	Expectancy-value-principle	Built in as part of problem solving
4	Common methods within connectionist modules	Could be implemented in rule modules	Specific to module type. Can vary	None built in (users can specify)
5	None built in (users can specify as required by their architecture)	None built in (users can specify as required by their architecture)	Includes BDI computation model	Includes BDI computation model
6	None built in (e.g., Wright et al. 1996, included simple forms of deliberative mechanisms and meta-management)	None built in (users can specify as required by their architecture). Logan's A* with bounded constraints available, among others	None built in (users can specify as required by their architecture)	None built in (users can specify as required by their architecture)
7	Usually not extant	Varies, some models do well	Usually good; decision making domain of these models	Varies, often implicit

Table 5.1: Comparison of Architectures (continued)

Architecture	Multitasking	
	Serial/Parallel	Resource Representation
1 EPAM	Serial processing; learning done in parallel	Limited STM capacity, limited perceptual and motor resources (uses time parameters)
2 SDM	Fully parallel recall process, serial recall of sequences	Architecture too low-level for representation to be explicit
3 PSI	System tries to fulfill different needs (i.e., water, energy, pain-avoidance, etc.); interrupts goal-directed behavior to profit from unexpected opportunities	Allocation of time to run intention according to strength of underlying need and according to expectancy of success
4 COGENT	Modules can work in parallel, but information passed between them serially	Would vary with the knowledge included in modules
5 JACK as an example of BDI architectures	Supports multiple computational threads handled safely within the JACK Kernel—achieving atomic reasoning steps	Agents have time perception. Time can be real or simulated (dilated, externally synchronized, etc.)
6 Sim_Agent Toolkit	Discrete event simulation technique, with rule sets within each agent time-sliced, as well as different agents being time-sliced	Allocation of cycles per time-slice can be made for each rule set, or for each agent. No built-in memory resource limits. Will differ for each architecture type created
7 Engineering-based models (e.g., APEX)	Varies, sometimes explicit models	Varies. Those that interact with simulations more advanced

Table 5.1: Comparison of Architectures (continued)

		Multiple Human Modeling	Implementation Platform
	Goal/Task Management		
1	Bottom up + 1 main goal per task simulated	Potential through multiple EPAM modules	Mac, PC (any system supporting Common Lisp). Graphic environment supported only for Macintosh
2	None	None	UNIX (easily ported)
3	There is a steady competition of different needs/motives to rule. Strongest will win and inhibit others	Potential through multiple PSI models with different "personalities" by varying starting parameters. Multiple agents can run in same environment, see each other, interact, and, to a certain degree, communicate	Windows 95, 98, 2000, NT
4	None built in. Users can specify through module selection and programming	None	UNIX (X windows). Microsoft Windows
5	Built in. JACK Language includes: wait_for (condition), maintenance conditions, meta-level reasoning, etc.	Allows multiple agents, running together or distributed, to interact and communicate as a team or as adversaries. Extensions to the basic model (e.g., team models also allowed)	Runs on all platforms that support JAVA 1.1.3 or later. Graphic components (i.e., development environment) require JAVA 2 v 1.2 or later
6	None built in (users can specify as required by the architecture)	Toolkit allows multiple agents to sense one another, act on one another, and communicate with one another	Runs on any system supporting Poplog (and for graphics, X window system). Tested on Sun/Solaris, PC/Linux, DEC Alpha/UNIX. Should also run on other UNIXes and VAX VMS. Should work without graphics on Windows NT Poplog
7	Varies. Some advanced	Some have none; some work in teams	Varies. Not usually designed for dissemination

Table 5.1 Comparison of Architectures (continued)

Architecture	Implementation Language	Support Environment
1 EPAM	Common Lisp	Lisp programming + editing tools. Some graphic utilities for displaying eye movements, structure of discrimination tree, and task. Customized code used for each task modeled
2 SDM	C, JAVA	None
3 PSI	Pascal (Delphi 4)	Delphi 4 features
4 COGENT	Prolog	Graphic and textual editors
5 JACK as an example of BDI architectures	JAVA. JACK written in and compiles into pure JAVA	JACK Make utilities, and all available JAVA tools. JACK development environment (JDE) provides GUI for creating and editing agent structures. Further debugging and visualization tools under development
6 Sim_Agent Toolkit	Pop-11 (but allows invocation of other Poplog languages (Prolog, Common Lisp, Standard ML, & external functions, e.g., C)	Poplog environment, including VED/XVED, libraries, incremental compiler, etc
7 Engineering-based models (e.g., APEX)	Varies	Often simple

Table 5.1: Comparison of Architectures (continued)

	Validation	Comments
1	Extensive at many levels	EPAM models focus on single, specific information processing task at a time. Not yet scaled up to multitasking situations. Used in high-knowledge domains (e.g., chess, with about 300,000 chunks)
2	None	SDM should be seen as system component (e.g., good way of representing long-term memory for patterns and motor behaviors in larger system)
3	Achievement data and parameters of behavior compared between subjects and models in two different scenarios (BioLab and Island). Different human subjects can be modeled by varying parameters	
4	Would be by architecture. Some have been done by modeling previously validated models	
5	Would be by architecture. None known	
6	Would be by architecture. None known	
7	By model. Usually validated with expert opinion. Some may be compared with data	Wide range of models here

CHAPTER 6

Review of Recent Developments and Objectives: Specific Projects

We now examine specific projects within the general application areas noted in Chapters 2, 3, and 4, broadly grouped into projects that support the objectives in the previous chapters, that is, of providing more complete performance, supporting integration of models, and improving model usability. The format of the projects follows the general format used in Pew and Mavor (1998). Where appropriate, this summary also comments on the feasibility and concerns that may arise if the projects are implemented in Soar, a current common approach for computer-generated forces. The estimates are uniformly optimistic to allow comparisons. The estimates are in terms of programmer or analyst time, and assume adequate supervision and cooperation with other organizations.

6.1 Projects Providing More Complete Performance

The projects presented here address the issues raised in Chapter 2. They are grouped into three main categories. We also note some potential additional uses for models of behavior as well as current uses in synthetic environments.

6.1.1 Gathering Data From Simulations

It is very clear and consonant with Pew and Mavor (1998, chap. 12) that data need to be gathered to validate models of human and organizational behavior. An approach at which they hint is to instrument synthetic environments. Synthetic environments should be instrumented not only for playback, but also in a way to provide data for developing and testing models. While the data are not directly equivalent to real-world behavior, as the environment becomes more realistic the data should become more realistic as well.

A uniform representation for data from simulations should be created. This representation should be readable by humans, at least in some formats.

Creating summary measures will also be necessary. Otherwise the sheer volume of data may preclude its analysis. The individual actions of control are not likely to be useful on their own (e.g., pressing an accelerator) but will be required to build higher-level summaries. Creating these summaries is likely to represent an additional research agenda item requiring AI, domain knowledge, and some understanding of behavioral data.

The playback could be quite large for developing models. Analysis of data from synthetic environments might also provide insights into the quality of the simulation (e.g., how quickly someone could act and whether they were limited by the simulation's ability to display information) and provide insights about the implementation of doctrine (e.g., how often tank crews actually follow doctrine). When done in cooperation with a simulator's

developers, the resources required for this task could be quite modest. Otherwise, it could take some time. Developing initial automated summaries is a 6- to 12-month effort.

6.1.2 Understanding Expectations of Behavior

Providing realistic behavior requires understanding what people expect from other people and what aspects of an adversary are necessary for training. (These two may be quite different.) In one sense, this means understanding the Turing test: what is necessary to appear human? More important, however, is knowing what is necessary to train people. A model that passes the Turing test and appears human might be weaker or unusual in some way. Thus, training with the model might not result in transfer of learning or result in learning an incorrect behavior.

A useful exercise would be to study which characteristics of behavior make a model appear human (so-called believable agents). The model must start with competencies; it must be able to perform tasks. It should also include errors, hesitations, and variations in behavior.

Work with the Soar Quake-bot on how firing accuracy and movement speed make agents believable is an example of what is required (Laird & Duchi, 2000) to understand what people think is human. The Soar Quake-bot has been evaluated on such things as firing accuracy with observers asked to rate its humanness. The measure of humanness, however, does not reveal how good the Quake-bot is with respect to training. Nor does it reveal what aspects of the Quake-bot should be made more (or less) human to improve training. The current belief is that appearing (or behaving) more human makes a better opponent to train against, but we do not know of any evidence to support this belief.

Another example is the Fuzzy Logic Adaptive Model of Emotions (FLAME). In this work (Seif El-Nasr, Yen, & Ioerger, 2000), several models of a pet that followed the user around in a virtual house were tested for believability. The model that included learning and fuzzy behavior was the judged the best. While not a complete test, this type of project starts to find out what makes agents believable through tests. In this example, learning and emotions were both helpful.

A useful 6-month to 1-year study would be to examine a range of models and humans in a synthetic environment, noting observers' comments and behavior toward a range of behavior. It might be that these aspects make an agent appear human, but it might also include implicit effects, such as second-order (or lagged) dependencies in behavior. The results would be important for training and also useful for creating models used in analysis. This project is similar to, but conceived of separately from, a similar call proposed by Chandrasekaran and Josephson (1999) to develop a better understanding of how to and how far to develop models.

The results are also essential for understanding how to help modelers. The results will point out the most likely mismatches to be left in models because modelers do not consider such behavior abnormal. This will provide suggestions about where comparisons with data are particularly needed by models. As this is basically experimental work, less resources are needed, but the time to run the experiment and analyze the results will take up to a year for preliminary studies.

6.1.3 Including Learning in Models

Work on creating agents in synthetic environments has been successful, however, one particularly useful aspect that has not been modeled is learning. A worthwhile project would be to take a learning algorithm and put it to use within a synthetic environment, either as part of a problem solver or as an observer. There are a variety of learning algorithms and models that would be appropriate. Some examples include: connectionism; one of the hybrid learning architectures developed within the ONR program (Gigley & Chipman, 1999); Programmable User Models (Young, Green, & Simon, 1989b); Soar with learning turned on; ACT-R; EPAM; or any of a wide variety of machine-learning algorithms.

Creating a model that learns will be difficult. This task is large and would allow multiple subprojects to be attempted. It could be supported by a wide range of resources. Including learning with problem solving has been difficult in the past, but it is likely to lead to more accurate agents that may be useful for testing and developing tactics.

Soar models exist that function fairly well in a synthetic environment. If these could be used, a small project of a programmer-year or two should be able to create an initial model that learns in a synthetic environment. Attaching a learning component to find regularities in behavior is likely to take at least that much time. Both projects would provide potential PhD topics and are broad enough to be supported by a wide range of resources.

6.1.4 Including a Unified Theory of Emotions

There are three specific projects related to modeling emotions and other behavioral moderators in architectures that we can propose: (1) adding general emotional effects, (2) adding reactive emotions, and (3) testing emotional models with performance data. While work is ongoing implementing models like this in Soar (Chong, 1999; Gratch, 1999) and ACT-R (Belavkin et al., 1999), the domain is large. Projects can range from a few months to implement a simple emotional effect to several years or decades to incorporate a significant amount.

Adding general emotional affects. As noted above, it is possible to start to realize emotions and affective behavior within toolkits like Sim_Agent and general cognitive architectures like ACT-R and Soar. Including emotions will provide a more complete architecture for modeling behavior and a platform for performing future studies of how emotions affect problem solving. Including emotions may also provide a way to duplicate personality and provide another approach to account for appropriate variations in behavior. Hudlicka (1997) provides a list of intrinsic and extrinsic behavior moderators (similar to the categories suggested in Ritter, 1993b) that could be modeled. Boff and Lincoln (1988) provide a list of regularities related to fatigue and other related stressors that might be considered for testing against a general model of emotional effects. For example, by making ACT-R's motivation sensitive to local performance (rule successes and failures), we have fit the Yerkes-Dodson law (Belavkin & Ritter, 2000).

These models should move from applying to a single task to multiple tasks. They would then become modifications to the architecture and thus reusable.

Adding reactive emotions. Modeling reactive and long-term moderators as well as slower-acting behavioral moderators is worthwhile. The effect of stress also changes the state of the competence in important ways. A proportion of troops engaged in active combat will become ineffective as a result of fear and stress-fatigue. Stress would also be increased by the number of casualties taken by a given platoon, length of time without sleep, weather conditions, perceived chance of survival, and so on. Modeling these effects at the micro-level of individuals, following known distributions, would advance the realism of simulations in interesting ways and support teaching existing doctrine.

In production system architectures, these emotions can initially be implemented by changing the decision (rule-matching) procedure, adding rules to make parameter changes, and by augmenting working memory to include affective information (e.g., an operator or state looks good or bad). These types of changes are being applied to an existing model, which matches adult behavior well, to better match children's more emotional behavior (Belavkin et al., 1999). These emotional effects should improve the match to the children's performance by (1) slowing down performance in general, (2) slowing down initial performance as the child explores the puzzle driven by curiosity, and (3) abandoning the task if performance is not successful. This work should be extended and applied more widely.

Testing emotional models with performance data. Many of the theories of emotions proposed have not been compared with detailed data. Partly this may be because there is not always a lot of data available on how behavior changes with emotions. It is no doubt a difficult factor to manipulate safely and reliably. But the models must not just be based on intuitions.

The use of simulators may provide a way to obtain further data with some validity. Better instrumentation of some primary features of emotions (e.g., heart rate, blood pressure) is providing new insights (Picard, 1997; Stern, Ray, & Quigley, 2001) and will be necessary for testing models of emotions.

Some argue that emotions are necessary for problem solving. Examples of brain-damaged patients (e.g., Damasio's Elliot [1994]), who have impaired problem solving and impaired emotions, are put forward. It is not clear that emotions per se are required, or if multiple aspects of behavior are impaired as well as emotions by the trauma. Others argue from first principles that emotions (realized as changes in motivation due to local success and failure during problem solving in this example) can improve performance (Belavkin, 2001). A model compared with data may help answer whether this is true. Clearly, AI models of scheduling do not have the same troubles scheduling an appointment despite their lack of emotion.

6.1.5 Including Errors

There are two premises that underpin the modeling of erroneous behavior. The first premise is that the attribution of the label *error* to a particular action is a judgment made in hindsight. The identification of the erroneous action forms the starting point for further investigation to identify the underlying reasons *why* a particular person executed that particular action in that particular situation. In other words, the erroneous action arises as the result of a combination of factors from a triad of sources: the person (psychological and

physiological factors), the system (in the most general sense of the term), and the environment (including the organization in which the system is deployed).

The second premise acknowledges that an error is simply another aspect of behavior. In other words, any theory of behavior should naturally encompass erroneous behavior. The behavior can be judged as erroneous only with respect to a description of what constitutes correct behavior.

Once these premises are accepted, it becomes apparent that modeling erroneous behavior is actually an inherent and important part of modeling behavior. If the psychological and physiological limitations of human behavior are incorporated into a model of human behavior, then particular types of erroneous behavior should naturally occur in certain specific situations. The corollary of this argument is that an understanding of erroneous behavior can be used as the basis for evaluating models of behavior. So, if a human performs a task correctly in a given situation, the model should also be able to perform the task correctly in the same situation. If the situation is changed, however, and the human generates erroneous behavior as a result, the model should also generate the same erroneous behavior as the human in the new situation, without any modifications being required to the model.

Modeling error therefore depends on understanding the concept of error—its nature, origins, and causes—and central to this is the need for an accepted means of describing the phenomenon (Senders & Moray, 1991). In other words, a taxonomy of human error is required with respect to these tasks.

The utility of the taxonomic approach, however, depends on the understanding that the taxonomy is generated with a particular purpose in mind. In other words, the taxonomy has to reflect:

- A particular notion of what constitutes an error.
- A particular level of abstraction at which behavior is judged to be erroneous.
- A particular task or domain.

There is a need to be very clear about the classes of errors and their origin in the models so that the appropriate ones can be included. In the military context, for example, a major source of error is communication breakdown. One approach to developing an appropriate taxonomy of errors for the military domain is to use the scheme that lies at the heart of the CREAM (Hollnagel, 1998). The CREAM purports to be a general purpose way of analyzing human behavior in both a retrospective and a predictive manner. Although the method was developed on the basis of several years of research into human performance, mainly in the process industries, it is intended to be applicable to any domain.

The CREAM uses a domain-independent definition of what constitutes an erroneous action (also called error modes or phenotypes). One of the goals of the CREAM is to be able to identify the chain of precursors for the various error modes. Identifying the chain is achieved by means of a set of tables that define categories of actions or events. At the highest level, there are three types of tables:

- Human (or operator),
- Technological (or system), and
- Organizational (or environment).

Within these categories there are sub-category tables. So, for example, the human tables include observation, interpretation, planning, and so on.

The individual actions or events are paired together across tables on the basis of causality or, to use a more neutral term, in a *consequent-antecedent relationship*. When the CREAM is used to analyze a particular accident or incident retrospectively, the aim is to build up the list of possible chains of events and actions that led to the accident or incident.

The contents of the tables are domain-specific, so the first step in developing the taxonomy for agents in synthetic environments depends on identifying the appropriate categories of events and actions for the military domain. These categories and the links between individual actions or events will be generated from a combination of knowledge elicited from domain experts and a review of the appropriate literature.

The second step is to generate the possible chains of actions and events that precede the various error modes, based on information available from reports of real accidents or incidents. This process will involve access to desensitized accident or incident reports—like those used in the Confidential Human Factors Incident Reporting Programme (CHIRP; Green, 1990) originally operated by the RAF Institute of Aviation Medicine—that can be analyzed and coded using the domain-specific CREAM tables generated in Step 1. Where omissions from the tables are detected, or links between actions do not already exist, these should be added to the tables.

The possible causal chains of events or actions generated by the second step will provide the basis for a specification of behavior in a particular situation. Models of behavior should yield the same sequences of actions and events in the same circumstances. The specification of behavior can therefore be used to test the models of behavior for compliance, during development, with the model being modified as appropriate to match the specification.

In addition, the results of the analysis of the incident behavior provide a basis for evaluating the veracity of synthetic environments. Performance in the real world (as described in the incident reports) can be compared with the way people behave when performing in the synthetic environment. There should be a high degree of correspondence between the two. If there is a mismatch, the mismatch suggests that there is a difference between the real world and the synthetic environment, which may be worth further investigation to identify the source of the difference.

One other beneficial side effect of the CREAM analysis is that the resultant chains of actions and events can be used in training personnel to manage error. If common chains of actions or events can be identified, it may be possible to train personnel to recognize these chains, and take appropriate remedial action before the erroneous action that gives rise to the incident is generated.

Initial models that include erroneous behavior can best be created with an existing model. One to three years of work should lead to an initial model that includes some errors and has been validated against human behavior.

6.1.6 Including a Unified Theory of Personality

It would be useful to identify features that lead to modeling personality, problem-solving styles, and operator traits. While models that choose between strategies have been created, there are few models that exhibit personality by choosing between similar strategies (although see Nerb, Spada, & Ernst, 1997, for an example used to put subjects in a veridical, but artificial, social environment). Personality will be an important aspect of variation in behavior between agents.

Including personality requires a task (and the model) to include multiple approaches and multiple successful styles. It is these choices that can thus appear as a personality. If the task requires a specific, single approach, it is not possible to express a strategy. Psychology, or at least cognitive psychology, has typically not studied tasks that allow, or particularly highlight, multiple strategies. Looking for multiple strategies has also been difficult because it requires additional subjects and data analysis that before has not represented real differences in task performance. Differences in strategies, however, lead to variance in behavior (e.g., Delaney, Reder, Staszewski, & Ritter, 1998; Siegler, 1987).

There appear to be at least the following ways to realize variance in behavior that might appear like personality: learning, differences in knowledge, differences in utility theory and initial weighting, and differences in emotional effects. Including a subset of these effects in a model would fulfill a need for a source of regular, repeatable differences between agents in a situation.

All of the current cognitive architectures reviewed here and in Pew and Mavor (1998) can support models of personality. These types of changes should be straightforward, as long as there are multiple strategies. In Soar, personality can be expressed as differences in task knowledge, as well as differences in knowledge about strategy preferences either absolutely or based on different sets of state and strategy features. ACT-R appears to learn better and faster which strategy to use compared with a simple Soar model, but ACT-R requires additional state features (Ritter & Wallach, 1998). Models in both architectures can, however, modify their choice of strategies. The role of (multiple) strategies has been investigated within the EPAM architecture in several tasks, such as concept formation (Gobet et al., 1997) and expert memory (Gobet & Simon, 2000; Richman, Gobet, Staszewski, & Simon, 1996; Richman et al., 1995).

These models could also be crossed with emotional and other non-cognitive effects to see how personality types respond differently in different circumstances (broadly defined). This could even be extended to look at how teams with different mixes of personalities work together under stress.

The amount of work to realize a model in this area will depend on the number of factors taken into account by the model. Providing a full model of personality and how it interacts with tasks and with other models is a fantasy at this point. However, a minimal piece of work would take an existing model and give it more of a personality. A more extensive

project over a year or two would apply several of these techniques and see how it starts to match human data.

6.1.7 Including a Model of Situation Awareness and Rapid Decision Making

Novices have to do problem solving. Experts can do problem solving but save effort (or improve their problem-solving performance) by recognizing solutions based on the problem. Viewed broadly, a model that does this transition starts to provide an explanation of situation awareness and Rapid Decision Making (RDM) as a result of expertise and recognition.

Able (Larkin, 1981) and its recent re-implementations (Ritter et al., 1998b) provide a simulation explaining the path of development from novice to expert in formal domains (i.e., those where behavior is based on manipulated equations such as physics or math). The early (or barely) Able model works with a backward-chaining approach, that is, it starts with what is desired and chooses domain principles to apply based on what will provide the desired output. This approach is applied recursively until initial conditions are found. The chunking mechanism in Soar gives rise to new rules that allow the model to use a forward-chaining method that is faster. That is, from the initial conditions new results are proposed. The rules are applied until the desired result is found. Students at the University of Nottingham have applied the Able mechanism to several new domains. Their examples are available at www.nottingham.ac.uk/pub/soar/nottingham/student-projects.html.

Work could be done to translate this mechanism, which has worked in Lisp and in several versions of Soar, into other architectures and extend it from a simulation to a full process model. This would require a rather modest amount of effort, less than a programmer-year to get started if the programmer was familiar with Soar. Applying it in a realistic domain would take longer.

6.1.8 Using Tabu Search to Model Behavior

The internal architecture of a combatant might be constructed from a perceptual module that is closely coupled to the synthetic environment and can be modified by plug-in items that alter the incoming data to be processed (night-vision aids, etc.). The results of perception are crudely classified using a learning system such as a multi-layer perceptron, which triggers a rapid emotional response and consequent reactive behavior. This behavior might be generated using an SDM that finds the nearest match to previous scenarios and is capable of producing a sequence of outputs rather than a single-state result. Both perception and emotional response are calibrated by a perceptual and personality model that may be unique to individual entities, albeit assigned from a known distribution.

The cognitive processing would be rule-based using an established cognitive model, for example, ACT-R, with planning activities augmented by a Tabu search. There would be interactions between the state of the entity (including its emotional state) and the cognitive processing based on psychological data on human performance under stress. This approach is similar and perhaps a generalization of Sloman's meta-architecture, and the Soar and PSI architectures.

6.2 Projects Supporting Integration

The projects presented here roughly address the issues raised in Chapter 3. Integration is approached in two ways here: integrating model components and integrating the model with simulations in more psychologically plausible ways. Several projects described in this subsection could be equally at home in the set of projects for making modeling routine because the two areas are related.

6.2.1 Models of Higher-Level Vision

It has been argued that an understanding of higher-level vision is necessary for continued development of models in synthetic environments (Laird, Coulter, Jones, Kenny, Koss, & Nielsen, 1997) and we agree (Ritter et al., 2000). Neisser's (1976) perceptual cycle is just starting to be explored with models.

There are several areas of Higher-Level Vision (HLV) that are of particular interest for military modeling. These areas include:

- How information from long-term memory indicates incoming danger or serious change in the environment.
- How HLV directs attention.
- How HLV integrates various aspects of information, or integrates information occurring at different times.
- How HLV can be used to facilitate learning.
- How HLV can be used in planning and problem solving.

To put it simply, HLV is at the interface between Lower-Level Vision (LLV) and postulated memory entities such as productions, schemas, concepts, and so on. At the present time, this interface is poorly understood, perhaps because LLV and long-term memory are not understood in a sufficiently stable way. (However, see Kosslyn & Koenig, 1992, for neuropsychological hypotheses about HLV.)

Most models of cognition such as Soar and ACT-R (actually, most architectures reviewed by Pew & Mavor, 1998) use modeler-coded information, which avoids dealing with the interface between LLV and long-term memory constructs. Neural nets for vision have been used to go from pixel-like information to features or even higher but have not been incorporated into higher-cognition models. CAMERA (Tabachneck-Schijf, Leonardo, & Simon, 1997), and to a certain extent EPAM (Feigenbaum & Simon, 1984; Richman & Simon, 1989), explore ways in which features may be extracted from low-level representation, and may be combined into long-term memory constructs.

The relationship of HLV and problem solving is undoubtedly an area where more research should be carried out. For example, modeling instruction and training requires a theory of how low-level acoustic input merges with low-level visual input and connects to long-term memory knowledge. In some cases vehicles and gunfire will be heard rather than seen and sounds will direct visual attention in the appropriate direction. Perceptual models of hearing are also well-developed and exploited with dramatic success in, for example, the

MPEG-2 compression standard that is likely to form the basis for much broadcast and recorded sound in the future. The variance among individuals is large for both auditory and visual perception, and both processes are degraded temporarily or permanently by intense overload, as is likely in a military environment.

Work extending this approach to create integrated architectures (Byrne et al., 1999; Hill, 1999; Ritter & Young, 2001) is ongoing. Significant progress will require at least a year-long project, and a longer period would be more appropriate.

6.2.2 Tying Models to Task Environments

Tying cognitive models to synthetic environments in psychologically plausible ways should be easier. There are two approaches that seem particularly useful and plausible that we can ground with particular suggestions for work. They are consistent with Pew and Mavor's (1998, p. 200) short-term goal for perceptual front-ends.

The first approach is to provide a system for cognitive models to access ModSAF's display and pass commands to it. This approach has the advantage that it hides changes in ModSAF from the programmer/analyst and from the model. The disadvantage is the need for ModSAF experts, programmers, users, time, and money to make it work. There has been such a system created for Soar models to use ModSAF (Schwamb, Koss, & Keirse, 1994), but it is our impression that this system, although it was quite useful, needs further development and dissemination.

The second approach is to create a reusable functional model of interaction based on a particular graphics system or interface tool (as does the Nottingham Functional Interaction Architecture and ACT-R/PM). A functional rather than a complete model may be more appropriate here as a first step. This functional approach has been already created in Tcl/Tk (Lonsdale & Ritter, 2000), Garnet and Common Lisp (Ritter et al., 2000), Visual Basic (Ritter, 2000), Windows bitmaps (St. Amant & Riedl, 2001), Windows 98 objects (Misker, Taatgen, & Aasman, 2001), and most recently in JAVA. They could be created in Amulet, X-windows, Delphi, or a variety of similar systems, each of which allows models to interact with synthetic environments through a better programming interface. A functional model would then provide the necessary basis for improving the accuracy and psychological plausibility of interaction.

This approach to providing models access to information in simulations could also support creating cognitive models in general, such as for problem solving, working memory, and the effect of visual interaction. These could be later assimilated back into models and architectures in the synthetic environments.

An excellent programmer very familiar with their language can now create an initial system in about two weeks. Integrating and applying these models takes several months to a year.

6.2.3 Ongoing Review of Existing Simulations

To provide for reuse and to understand the current situation, a review of simulation systems used (for as broad a geographic region as possible, working with allied nations if possible) should be created that is similar to the listing in Pew and Mavor (1998, chap. 2

Annex). This listing, for example, could initially be created by an intercalated year (co-op) student and then maintained as part of standing infrastructure. This listing could provide an initial basis for understanding what the total needs were and the totality of current simulation efforts. While the U.S. Defense Modeling and Simulation Office may do this in the United States, we do not know of similar efforts in the United Kingdom.

6.2.4 Focus on a Flagship Task

Supporting all the uses of synthetic forces as shown earlier in Table 1.1 with a single model of behavior is probably impossible in the short term. The uses of simulations in operations research, training individual group behavior, and examining new materials or doctrine are too disparate to be met by a single approach. While the various levels and uses of simulations mentioned here are related by the real world they all represent, it does not appear to be possible in the next 5 to 10 years to integrate them to the extent to which the real world is integrated.

While there may be some systems that allow multiple use, and there will certainly be some reuse between these areas, a focus for work must be selected. Therefore, a more narrow focus on the most important uses should be adopted by funding agencies. Taking a more focused approach appears to be happening in several places already. A selective focus on the most approachable or natural set of uses is more likely to be successful in the short term and may provide a better foundation upon which to build in the long term. Discussion of these issues should be grounded, if possible, with a set of potential uses with possible systems and domains that will be used in the next 5 to 10 years. Complete unification is not likely in that time period, nevertheless significant reuse should be sought.

Having a focus would also support the choice of a specific application. Applications can then be chosen with a user audience in mind. Having a specific audience will help the application to be useful and seen as useful by a well-defined user community.

Work that attempts to serve too many needs will serve all of them poorly. Projects and research programs will have to pick a domain and an application (or two), and work with them. This application could be an existing use or application or it could be a new use. Work with simulations for training often have high payoffs. Augmenting existing training would be a natural place to consider starting.

The students being trained could also be used to help test the simulation. Apocryphal tales from MIT suggest that building computer-based tutors to deliver instruction is as useful for learning as using the resulting tutors. Creating and validating these models would be good training for such students as well.

6.2.5 A Framework for Integrating Models With Simulations

Perhaps the most significant current requirement is a way to integrate multiple cognitive and behavioral architectures into synthetic environments. Currently, it takes a large amount of effort to introduce new models of behavior and connect them directly to simulations via the Distributed Interactive Simulation (DIS) protocol. Coupling cognitive architectures to a simulation via ModSAF is probably marginally easier because ModSAF, while difficult to use, provides physical models and an interface to the network. The left-hand side of Figure

6.1 shows the organization of systems like Tac-Air Soar that interact with ModSAF to generate behavior.

A worthwhile medium- to long-range goal would be to develop utilities to support making a tool like ModSAF even more modular. The core activities of supporting communication across the network for simulation and supporting the physical model need to be provided, but are not of particular interest for modeling behavior.

Efforts have attempted to provide similar interfaces for Soar; however, they have never been fully successful. They have made hooking up Soar easier but have not yet made it easy (e.g., Ong, 1995; Ong & Ritter, 1995; but also see the most recent work by Jones, 2001, and Wallace, 2001). Work on the Tank-Soar simulator (provided as a demo in the latest release of Soar, Soar 8.3) might provide a path for this.

The right-hand side of Figure 6.1 shows how future systems might interact with ModSAF using the same interface that users see through a simulated eye and hand designed to allow models to interact with synthetic environments (Ritter, Jones, Baxter, & Young, 1998a). The interface to the physical simulation could no doubt be made more regular and easier to use so that other architectures, such as Sim_Agent, could be hooked up to it. We suspect this project might take a good programmer familiar with ModSAF about half-time over a year because we had a similar system built in 2 weeks by someone who was an expert in their graphic programming language. A much longer time should be allowed. This system requires knowing ModSAF very well because it will make use of all of ModSAF and may require extending ModSAF.




Figure 6.1: On the left, a functional description of Tac-Air Soar and how it uses ModSAF. On the right, a perceptual interface to ModSAF.

6.2.6 A Framework for Integrating Knowledge

Currently, there are multiple knowledge sets (as models) in different simulations that exist in multiple formats. It would be useful to create a framework for integrating multiple knowledge sets, allowing the knowledge to be reused in different simulations.

One way to create a framework for integrating knowledge is to create a task editor that could take a knowledge set and compile it for different architectures. The editor would have to be based on a high-level description of knowledge, such as generic tasks (Wielinga, Schreiber, & Breuker, 1992). These generic tasks would then be compiled into things such as an ACT-R or Soar rule-set.

There are potentially huge payoffs from this very high-risk project. First, this project would provide a way to reuse knowledge in multiple simulations. Second, the reuse that would arise would help validate models and might provide a way forward for validating architectures. Third, this project would provide another way of documenting behavior models. The (presumably) graphic representation would allow others to browse and understand the model on a high level. Fourth, it would assist in writing models. In most cases, there are a lot of low-level details in creating these models that are not of theoretical interest but require attention, such as using the same attribute name consistently (recent Soar interfaces now support this). A high-level compiler for knowledge like this would bring with it all the advantages traditionally associated with high-level languages. When done for Soar, the higher-level language allowed models to be built two to three times faster (Yost, 1992, 1993).

PC-Pack (www.epistemics.co.uk) is a potential tool to start building upon. Implementing an initial, demonstration version of this approach would take a good programmer 6 to 12 months. Putting it to use would take longer.

6.2.7 Methods for Comparing Modeling Approaches

We find ourselves in a position where a number of different approaches to simulating human behavior are available. Some of these approaches, at least, are based on datasets close enough to see themselves as rivals, and make competing claims about their suitability and quality. How can we assess and compare them?

There can, of course, be no one method that answers such a question. Earlier chapters of this report have discussed how practical considerations such as usability and communicability of models come into play as well as scientific qualities such as agreement with data. Thus, a wide range of comments about a model or architecture can be relevant to choosing between them.

However, there are some methods available that are too loose and varied to constitute a “technique” but are useful nonetheless for comparing and contrasting such differing approaches. They take the form of *matrix exercises*, in which a range of modeling approaches are pitted against a battery of concrete scenarios to be modeled. Young and Barnard (1987) provide the basic rationale for such a method and explain how it can be used to judge the fit and scope of a modeling approach. They argue, first, that the modeling approaches need to be applied to concrete scenarios. It is not sufficient to try comparing approaches on the basis of their “features” or “characteristics.” Second, it is important to use

a *range* of scenarios. Taking just a single case will inevitably introduce a bias towards or against certain approaches, and will fail to provide an indication of their scope. Young, Barnard, Simon, and Whittington (1989a) provide a short example of such a matrix exercise, and show how the entries in the matrix can be interpreted.

This kind of matrix exercise derives from the idea of a “bake-off” between rival approaches but also differs in important respects. There is unlikely to be a “winner,” one approach that is regarded as the best in all respects. Moreover, the matrix exercise is fundamentally cooperative rather than competitive. Instead of finding the “best” approach, bake-offs provide a tool for probing the scope of applicability of the different approaches, and investigating their relative strengths and weaknesses, advantages and disadvantages, for later modification and fusion. Pew and Mavor appear to call for this kind of activity (1998, pp. 336-339) as well.

Some exercises of this kind have been performed in public. At the Research Symposia associated with the CHI conferences in 1993 and 1994, Young (in 1993) and Young and C. Lewis (in 1994) organized such matrix exercises on the design of an *undo* facility for a shared editor (1993), and on the analysis of the persistent *unselected window* error and of the design of an automated bank-teller machine as a walk-up-and-use device (in 1994). Furthermore, there are precedents for such an exercise in a military research context. In 1993, NASA funded a comparative study of models of pilot checklist completion. The Office of Naval Research has funded, on a longer time scale, multiple analyses and modeling of several interactive tasks using hybrid architectures (Gigley & Chipman, 1999). The speech recognition community in the United States uses this approach in a quite competitive way as well.

The U.S. Air Force has recently started a similar program called Agent-Based Modeling and Behavior Representation (AMBR) to explore models of complex behavior (www.williams.af.mil/html/ambr.html). This multi-team project comparing four cognitive architectures was recently reported at the 2001 Computer Generated Forces Conference. For an overview, see Gluck and Pew (2001a; 2001b); Tenney and Spector (2001) provide a summary of the model to data fits in the most recent comparison round. Several more iterations of comparisons across architectures using different types of tasks are planned.

A final but important point about such an exercise is that it cannot be done successfully inexpensively. The exercise requires earmarked and realistic funding to provide useful results. A considerable amount of work is required: first in negotiating, agreeing, and then specifying a set of concrete and clearly described scenarios, ideally with associated empirical data; and then subsequently for applying the modeling approaches to the scenarios, performing the comparisons, and drawing conclusions. Multiple research groups are used, and the funding has been leveraged by the groups’ existing work and multiple funding sources.

6.2.8 (Re)Implementing the Battlefield Simulation

There are strong arguments for implementing communicating agents and intra-agent processes in JAVA. These are discussed in Bigus and Bigus (1997), and in the context of JACK Intelligent Agents by Busetta et al. (1995b). In fact, there are powerful arguments for building the entire synthetic agent simulation in JAVA as described below. This is possible

within the Higher-Level Architecture (HLA) framework. Implementing Soar in JAVA has also been mooted (Schwamb, 1998), as well as ACT-R (see www.jactr.sourceforge.net for information on a preliminary JAVA implementation of ACT-R 4, as of May 2001), although the usability of these architectures would suffer for this.

A core system implementation is needed that can then be accessed through Application Programming Interfaces (APIs). Supporting software is available for this, but any software could be developed for the purpose, provided it conformed with the standard. The core system might be written in JAVA or any other language provided only that an API is implemented. Similarly, entities may be written in any language, or several, provided that they set up calls to the API specification. There are a number of arguments for using JAVA as the basis for both individual entity simulation and for building a core system to the HLA specification. These are described below.

The single most attractive advantage of developing a synthetic battlefield simulation within a JAVA environment lies in the capabilities available within a Remote Method Invocation (RMI) that forms part of the JAVA run-time environment. This is a distributed object model with some similarities to Microsoft's Distributed Component Model (DCOM)[®] but with the advantage that it is effective on any platform that supports a JAVA run-time environment. It goes well beyond traditional remote procedure calls being entirely object-based, even allowing objects to be passed as parameters. Object behavior as well as data can be passed to a remote object in a seamless and transparent way. A mortar weapon being passed as an argument to an individual infantry man entity and arriving complete with its complement of munitions and ability to be fired gives a picture of this capability. The JAVA run-time environment also supports a naming and directory service API (JAVA JNDI) that allows the objects of RMI calls to be found. (For more details of this see www.javasoft.com/products/jndi/index.html.)

To show how such a service might be used, suppose that a simulation of an individual paratrooper has been developed. This simulation is a uniquely named JAVA object that can be invoked on any machine on the network used for the simulation. The JAVA Naming and Directory Interface (JNDI) service will inform a process about which machines have a suitable simulation available. To take this an important stage further, we use a class-factory object to produce the individual paratrooper objects. This class factory might use randomized parameters to make each entity distinct but fitting a known distribution (like Cabbage-Patch Dolls[®]). To introduce these entities into the simulation, a process would ask the naming service for a suitable class-factory object. This might be on one of any number of machines and is therefore extremely robust against damage to the network. The class factory can then be asked to produce any number of paratrooper entities, each of which (in JAVA) is capable of serializing itself to any other machine on the network, and running there. Indeed, the simulation can be moved from machine to machine at will, perhaps in response to a condition such as imminent power failure.

This approach would also support testing new platforms. A manufacturer might develop an improved simulation of a Tornado fighter-bomber. They then could introduce a new machine with a suitably registered class-factory object. Once this was connected to the network, the new simulation would be immediately available even if this were done while a simulation was running. No relinking, recompilation, or even pause in the simulation would be needed. The objects could be defined in conformity with the HLA standard.

JAVA also supports secure communications and has well-developed APIs for database connectivity and for driving graphics devices. An attractive user interface is very much easier to develop using the JAVA Foundation Classes (JFC) than, for example, using X-Motif. In addition, if a Just In Time (JIT) compiler is available to the RTE, programs developed in JAVA show little performance degradation in comparison with C++.

A synthetic environment could be developed using facilities offered by the JAVA runtime environment and existing APIs that would come much closer than existing simulations in meeting the design goals of maintainability, versatility, and robustness. This approach would have to be agreed upon by multiple communities and requires a large amount of resources to be applied uniformly.

6.3 Projects Improving Usability

The projects presented here roughly address the issues raised in Chapter 4. This section reviews several possible projects for making model building more routine. For practical reasons, it is useful to make the model-building process more routine. It is also important for theoretical reasons. If the models cannot be created within a time commensurate with gathering data, the majority of the work will continue to be data gathering because theory development will be seen as too difficult.

6.3.1 Defining the Modeling Methodology

There is not yet a definitive approach or handbook for building models that can also be used for teaching and practicing modeling cognitive behavior. Newell and Simon's (1972) book is too long and mostly teaches by example. Ericsson and Simon's (1993) book on verbal protocol analysis has comments on how to create models; although useful, the comments are short. VanSomeren, Barnard, and Sandberg (1994) provide a useful text, although it is slightly short and some of the details of going from model to data are not specified (if indeed they can be). Baxter's (1997) report and Yost and Newell's (1989) article are useful examples of the process, but both are tied to a single architecture and not widely available. There are other useful papers worth noting, but they are short and not comprehensive (e.g., Kieras, 1985; Ritter & Larkin, 1994; Sun & Ling, 1998).

Rouse (1980) has also made an attempt at describing the modeling process. He identifies the following steps as forming an important part of the modeling process: (1) definition, (2) representation, (3) calculation, (4) experimentation, (5) comparison, and (6) iteration. Rouse mainly focuses on the representation and calculation aspects of modeling, particularly from an engineering point of view. He describes several methodologies, including control theory, queuing theory, and rule-based production systems. He also provides a short tutorial on several of these modeling methods together with practical examples of systems engineering models. The examples are taken from a wide variety of domains including aviation, air traffic control, and industrial process control. It is not a complete treatise on human behavior, but does provide suggestions for methods that may be useful in modeling certain aspects of human behavior.

Similar tutorials and methodological summaries should be created until they converge. The results will be useful to practitioners and those learning to model; the latter will be an

important audience as this field grows. The output is most likely to require a textbook. A year to several years of support would significantly help create this set of learning materials.

6.3.2 Individual Data Modeling: An Approach for Validating Models

What is the best way to make theoretical progress in the study of behavior? Is it to develop micro-theories that explain a small domain or to aim at a higher goal, and develop an overarching theory covering a large number of domains—a unified theory? Modern psychology, as a field, has tended to prefer micro-theories. Unified theories have regularly appeared in psychology—think of Piaget’s (1954) or Skinner’s (1957) theories—but it is generally admitted that such unified theories have failed to offer a rigorous and testable picture of the human mind. Given this relatively unsuccessful history, it was with interest that cognitive science observed Newell’s (1990; see also Newell, 1973) call for a revival of unified theories in psychology.

One of the reasons for the limited success of Newell’s own brand of UTC is that the methodology commonly used in psychology, based on controlling potentially confounding variables by using group data, is not the best way forward for developing UTCs. Instead, Gobet and Ritter (2000) propose an approach, which they call Individual Data Modeling (IDM), where (1) the problems related to group averages are alleviated by analyzing subjects individually on a large set of tasks, (2) there is a close interaction between theory building and experimentation, and (3) computer technology is used to routinely test versions of the theory on a wide range of data. They claim that there are significant advantages here, that this approach will also help traditional approaches progress, and that the main potential disadvantage—lack of generality—may be taken care of by adequate testing procedures.

IDM offers several particular advantages in this area. It does not require as much data because the data will not be averaged but compared on a fine-grained level. Not requiring a large amount of data is attractive when the data are detailed or expensive to acquire, or where the model makes detailed predictions. The other advantage is that it provides a model that produces more accurate behavior on a detailed level. It is this detailed level of behavior that will be necessary to not only allow a model to appear human in a Turing test, but also lead to accurate training results because it performs like a comparable colleague or foe.

Work using IDM is ongoing at the University of Nottingham and at Pennsylvania State University. A full test would require one to two years of work to gather data and compare it with a model. Developing the IDM methodology and applying it could be combined with other projects, however, because it is a methodology and not a feature of behavior to include in a model.

6.3.3 Using Genetic Algorithms to Fit Data

There are two potential uses of genetic algorithms worth highlighting. The first is for generating behavior as described above in Section 5.2.1. The second is for optimizing model-fits by adjusting their parameters (Ritter, 1991). Most model-fits have been optimized by hand, which leads to absolute and relative performance problems. In absolute terms, researchers may not be getting optimal performance from their models. In relative terms, comparisons of hand-optimized models may not be fair. (Sometimes even one model

is optimized and the other not.) In the case of models with multiple parameters (with submodels to include), this job is not tractable by hand.

The results obtained by optimizing models with genetic algorithms suggest that optimizations done by hand are likely to be inferior to those done by genetic algorithms (Ritter, 1991) or by other machine-learning techniques (Butler, 2000). Use of genetic algorithms (or similar techniques) would improve performance in absolute terms, provide fairer comparisons between models, and encourage the inclusion of parameter set behavior in model comparisons. Several years of a PhD student working within a project with a model to optimize is probably a good way to progress work in this area.

This optimization should initially be done with an existing model so that the developers of the interface have a ready-made model and audience. The basic approach is simple and robust, and should be straightforward to demonstrate. Making the optimization routine and portable are separate and more advanced steps, so this project could take almost any amount of resources, ranging from a month to several years.

6.3.4 Environments for Model Building and Reuse

There remains a need for better environments for creating models. Few modeling interfaces provide much support for the user to program at the problem-space level or even the knowledge level, although the COGENT interface is interesting as an example of usability.

Soar, in particular, needs a better interface. While there is now a modest interface, even the latest versions of the Soar interface (Kalus & Hirst, 1999; Laird, 1999; Ritter et al., 1998b) are not as advanced as many expert system shells and are just becoming as comprehensive as the previous, Lisp-based version (Ritter & Larkin, 1994). The Soar interface is, however, providing increasing amounts of support at the symbol level (Jones, Bauman, & Laird, 2001; Roytam, 2001) and higher, including model-specific displays (Jones, 1999b). TAQL (Yost & Newell, 1989) and Able (Ritter et al., 1998b) have been moderately successful, but modest attempts to create high-level tools in Soar, for example. Gratch's (1998) planning-level interface should be expanded and disseminated as a modeling interface. Knowledge acquisition tools and techniques (e.g., Cottam & Shadbolt, 1998; O'Hara & Shadbolt, 1998) might be particularly useful bases upon which to build.

Associated with this project would be general support for programming. This includes lists of frequently asked questions, tutorials, and generating models or model libraries designed for reuse. These libraries should either exist in each architecture or in the general task language developed in the previous task. These would serve as a type of default knowledge for use in other applications. We can already envision libraries of interaction knowledge (about how to push buttons and search menus), arithmetic, and simple optimization like the default knowledge in Soar.

Work on improving the modeling interfaces for each architecture should be incorporated as part of another modeling project so that the developers of the interface have a ready-made audience. There are multiple additions that would be useful and multiple approaches that could be explored, so this project could take almost any amount of resources, ranging from a month to several years.

6.3.5 Automatic Model Building

Most process models induced from protocols are created by hand. There has been some work to do this automatically or semi-automatically with machine-learning techniques. Semi-automatic model generation is done in the event-structure modeling domain (a sociological level of social events) by a program called Ethno (Heise, 1989; Heise & Lewis, 1991). Ethno iterates through a database of known events finding those without known precursors. It presents these to the modeler, querying for their precursors. As it runs it asks the modeler to create simple qualitative, non-variablized token-matching rules representing the event's causal relationships based on social and scientific processes. The result at the end of an analysis is a set of 10 to 20 rules that shape sociological behavior in that area. In a sense, the modeler is doing impasse-driven programming (i.e., what is the next precursor for an uncovered event not provided by an already existing rule?). After this step, or in place of it, the modeler can compare the model's predictions with a series of actions on a sociological level (a protocol in the formal sense of the word). The tool notes which actions could follow and queries the modeler based on these. Where mismatches occur, Ethno can present several possible fixes for configuration. Incorporating the model with the analysis tool in an integrated environment makes it more powerful. It would be a short extension to see the social events as cognitive events in a protocol.

Stronger methods for building models from a protocol are also available. Cirrus (VanLehn & Garlick, 1987) and ACM (Langley & Ohlsson, 1984) will induce decision trees for transitions between states that could be turned into production rules given a description of the problem space, including its elements and the coded actions in the protocol. Cirrus and ACM use a variant of the ID3 learning algorithm (Quinlan, 1983). (ID3 induces rules that describe relationships in data.)

These tools look like a useful way to refine process models. Why is automatic creation of process models not done more often? Perhaps it is because these tools do not create complete process models. They take a generalized version of an operator that must be specified as part of a process model. It could also be that finding the conditions of operators is not the difficult problem but that creating the initial process model and operators is. It could also be that it is harder to write process models that can be used by these machine learning algorithms. In any case, these methods should be explored further.

Diligent (Angros, 1998), Instructo-Soar (Huffman & Laird, 1995), and Observe-Soar (van Lent, 1999) are approaches to create models in Soar that learn how to perform new tasks by observing behavior and inferring problem-solving steps to duplicate them. Related models have been used in synthetic environments (Assanie & Laird, 1999; van Lent & Laird, 1999). They have had limited use but suggest that learning through observation may be a way to create models as it is an important way that humans learn. Their lack of use could simply be due to the fact that they are novel software systems. As novel systems they are probably difficult for people other than their developers to use and will have to go through several iterations of improvement (like most pieces of software) before they are ready for outsiders. With a small user base (so far), the need has not forced software development, which has further decreased their potential audience.

Automatic modeling tools need to be developed. Machine-learning algorithms and theories of cognition are developed enough that this could be a very fruitful approach. A several-year effort here could yield large benefits of more routine modeling.

6.3.6 Improvements to ModSAF

A major problem with ModSAF is usability. ModSAF is large and has a complicated syntax. Users report problems learning and using it. One way to improve its usability might be a better interface; better manuals and training aids might also be useful.

The approach used by models of behavior to interact with basic simulation capabilities such as ModSAF needs to be regularized. A fundamentally better approach might be possible. There exists an interface between ModSAF and Soar that partly provides a model eye and hand. This eye/hand could be improved to provide a more abstract interface to ModSAF, one that might be easier to use (Schwamb et al., 1994).

One thing we have repeatedly noted is that getting models to interact with simulations is more bearable when both are implemented within the same development environment. When they are not, work proceeds much more slowly (Ritter et al., 2000; Ritter & Major, 1995), requiring a mastery of both environments. The situation is exacerbated because the development and use of any communication facility tends to be an ill-defined problem with numerous wild subproblems (i.e., problems where the time to solution can be high and with a large variance, that is, not easily predicted). So, for example, although the ModSAF Tac-Air system (Tambe, Johnson, Jones, Koss, Laird, Rosenbloom, et al., 1995) appears as if it was developed using joint compilation techniques, it was probably difficult to use because it implements communication between ModSAF and the Tac-Air model using sockets. Although informal communication with researchers in the Soar and robotics communities suggest that the use of sockets may be becoming more routine, this has not always been the case.

6.4 Other Applications of Behavioral Models in Synthetic Environments

There are numerous ways that behavioral models could be applied outside the military domain. We will examine four of them here.

The most obvious additional application of the models arising from approaches proposed in this report is in the provision of automated support for system operators. This support can take the form of intelligent decision-support systems or embedded assistants that guide operator behavior. There are some existing applications, most notably the Pilot's Associate (Geddes, 1989), its derivative, Hazard Monitor (Greenberg, Small, Zenyah, & Skidmore, 1995), and CASSY (Wittig & Onken, 1992), all from the aviation domain. In the United Kingdom, the Future Organic Airborne early warning system is attempting to insert a knowledge-based system into the Osprey aircraft and radar simulation to assist users.

These assistants, because they have a model of what the user is likely to do next, should be able to assist the user: if not by performing the task, then by preparing materials or information, or by modifying the display to help distinguish between alternatives or make performing actions easier. In the past, such assistants have had only a limited ability to model users. With increased validity and accuracy, these models may become truly useful.

The second application is in education and training. The uses in education have been fairly well illustrated by Anderson's work with cognitive model-based tutors (Anderson, Corbett, Koedinger, & Pelletier, 1995). In training, behavioral models can be used to provide experts to emulate and the same knowledge can also be used to debrief students' performances (Ritter & Feurzeig, 1988). The knowledge can also be used to populate adversaries and colleagues in the same environment (Bloedorn & Downes-Martin, 1985).

Training needs exist outside the military in several domains where dynamic models are necessary. Mining, for example, is starting to use virtual reality to train simple tasks (Hollands, Denby, & Brooks, 1999). Virtual reality is already being used to train hazard-spotting, avoiding mine machinery as a pedestrian, and driving vehicles underground (Schofield & Denby, 1995). A web search on virtual reality and training will indicate a wide range of other areas of application as well.

The third application is in entertainment. This has been proposed for some time as an application. A recent report by the U.S. National Research Council (Computer Science and Telecommunications Board, 1997) suggests that it is possible to use synthetic environments and the behavioral models in them for entertainment. This is currently being done by the Institute for Creative Technologies at the University of Southern California.

The fourth application is in systems analysis. The behavioral models can be used to examine different system designs to measure errors, processing rates, or emergent strategies. To return to mining again, truck models in a simulation can be used to examine road layouts in mines (Williams, Schofield, & Denby, 1998).

6.5 Summary of Projects

We have laid out important objectives for models of behavior in synthetic environments in the important areas of providing more complete performance, increased integration of the models with each other and with synthetic environments, and improved usability of the models. A wide range of funding bodies may be interested in supporting these projects because most of these projects have both engineering and scientific results. They will not only improve engineering models of human behavior, but they will also improve our understanding of behavior and our general scientific ability to predict and model human behavior generally.

These proposals, taken as a whole, call for several broad and general research programs. They suggest several moderating variables that affect cognition, including emotions and behavioral moderators, personality, and interactions with the environment, which should be included in cognitive architectures. They argue for creating or moving towards a more uniform format for data and models and a more clearly defined approach for modeling. There are also several concrete suggestions for making modeling easier and more routine, including providing more usable modeling environments and supporting automatic model generation. Finally, we were able to suggest some further applications of models of behavior in synthetic environments.

APPENDIX A

List of Participants

Attendees at the workshop on Techniques for Modeling Human Performance in Synthetic Environments, presented at the University of Nottingham on March 17, 1999, and their affiliations at that time.

Belavkin, Roman. University of Nottingham, rvb@cs.nottingham.ac.uk
Croker, Stephen. University of Nottingham, sfc@psychology.nottingham.ac.uk
Elliman, David. University of Nottingham, dge@cs.nottingham.ac.uk
Gobet, Fernand. University of Nottingham, frg@psychology.nottingham.ac.uk
Greig, Ian. DERA, igreig@dera.gov.uk
Howes, Andrew. University of Cardiff, howesa@cardiff.ac.uk
Logan, Brian. University of Birmingham, b.logan@cs.bham.ac.uk
Lonsdale, Peter. University of Nottingham, lpxprl@psychology.nottingham.ac.uk
Page, Ian. DERA, ipage@dera.gov.uk
Ritter, Frank (Chair). University of Nottingham, Frank.Ritter@nottingham.ac.uk
Russell, Simon. DERA, russell@dera.gov.uk
Shadbolt, Nigel. University of Nottingham, nrs@psychology.nottingham.ac.uk
Sheppard, Colin. DERA, csheppard@dera.gov.uk
Sloman, Aaron. University of Birmingham, a.sloman@cs.bham.ac.uk
Widdowson, Marc. VEGA Group, marc.widdowson@lineone.net
Young, Richard. University of Hertfordshire, r.m.young@herts.ac.uk

APPENDIX B

Description of Soar and ACT-R

Soar and ACT-R are two of the most commonly used cognitive architectures. They can be seen as theories of cognition realized as sets of principles and constraints on cognitive processing, a cognitive architecture (Newell, 1990). They both provide a conceptual framework for creating models of how people perform tasks. They are thus similar to other unified theories in psychology, such as PSI and COGENT.

Both Soar and ACT-R are supported by a computer program that realizes those theories of cognition. There are debates as to whether and how the theory is different from the computer program, but it is fair to say that they are at least highly related. It is generally acknowledged that the program implements the theory and there are commitments in the program that must be made to create a running system that are not in the theory—places where the current theory does not say one thing or another.

As cognitive architectures, their designers intend them to model the full breadth and width of human behavior. Such cognitive architectures, including the ones discussed in this report, do so to a greater or lesser extent, usually with the areas covered increasing monotonically over time. This approach to modeling human cognition is explained in books by Newell (1990) and Anderson (Anderson, 1993; Anderson & Lebiere, 1998). These books also provide introductions of Soar and ACT-R.

Further information on both Soar and ACT-R are available from the references cited here, as well as the sources included in the bibliography at the end of this appendix. The sources in the bibliography were used to write this appendix, particularly Johnson (1997), Jones (1996a, 1996b), and Ritter (2001).

B.1 Background of Soar and ACT-R

Soar and ACT-R are each based on a set of different theoretical assumptions, reflecting, largely, their different conceptual origins. Soar was developed by combining three main elements: (1) the heuristic search approach of knowledge-lean and difficult tasks, (2) the procedural view of routine problem solving, and (3) a symbolic theory of bottom-up learning designed to produce the power law of learning (Laird, Rosenbloom, & Newell, 1986). However, many of the constraints on Soar's theoretical assumptions consist of general characteristics of intelligent agents, rather than detailed behavioral phenomena. Soar's outlook is more biased towards performance because it arose out of an AI-based tradition.

In contrast, ACT-R grew out of detailed phenomena from memory, learning, and problem solving (Anderson, 1983, 1990; Singley & Anderson, 1989). ACT-R is thus suited more for predicting slightly lower-level phenomena, and is slightly more suited for predicting reaction times more accurately, particularly for tasks under 10 seconds in duration. These differences are relative; both architectures have been used for both high-

and low-level models, with attention paid to both performance and time predictions. ACT-R's outlook is more biased towards predicting reaction-time means and distributions because it arose out of a more experimental psychology tradition.

B.2 Similarities Between Soar and ACT-R

Soar and ACT-R can be seen as similar in numerous ways. They both have two kinds of memory, declarative (facts) and procedural (rules), although they represent these items differently. Typical instantiations of them now have input provided through a model of perception and output buffered through a model of motor behavior (Byrne, 2001; Chong, 2001; Ritter et al., 2000).

Both Soar and ACT-R model behavior by reducing much of human behavior to problem solving. Soar does this rather explicitly, being based upon Newell's information processing theory of problem solving (Newell, 1968), whereas ACT-R merely implies it by being goal-directed.

In both architectures these memories are conceptually infinite, with no provision being made for the removal of any memory item in ACT-R (the Soar architecture does perform removal of declarative memory, which therefore can be seen as a type of short-term memory). Manipulation of declarative memory can be accomplished by adding new items or changing existing ones. For procedural memory, rules may only be added to both architectures.

The course of processing involves moving from an initial state to a specified goal state. ACT-R has only one possible goal state (Version 5), whereas Soar may have several of them arranged in a stack. Movement between the initial and goal states usually involves the creation of sub-goals to accomplish the various parts leading up to the satisfaction of the goal.

Both ACT-R and Soar maintain a goal hierarchy where each subsequent sub-goal becomes the focus of the system. In ACT-R, these must be satisfied in a serial manner and in the reverse of the order they appear in the hierarchy (which is not directly visible to both the model and the modeler). Soar generally proceeds in a serial way as well, but is capable of removing (or solving) intermediate sub-goals should the current problem solving resolve a sub-goal that is much higher in the goal hierarchy. This difference makes ACT-R potentially less reactive, although work is in progress to make ACT-R more reactive (Lebiere, 2001).

B.3 Differences Between Soar and ACT-R

There are also fundamental differences between the two architectures. Soar only moves between states through changing the state as part of a decision procedure, which rules can vote on but cannot directly cause. In Soar, when no more productions can fire, an operator is selected or a state is modified. This whole process is called a decision cycle. Where an operator cannot be selected (e.g., due to preferences for the set of operators conflicting each other or not being complete), a sub-goal is created with a goal to choose the next operator. Movement between states is done in ACT-R by firing productions, which may change the state and goal stack directly.

Soar allows multiple rules to fire in parallel. This may lead to impasses because the knowledge in the rules may suggest different operators, but problem solving is available to resolve this. In ACT-R, when the conditions of several productions are met, a conflict resolution mechanism selects the production that it estimates to have the highest gain.

Learning in Soar occurs only for production memory. New rules are created by the architecture whenever a sub-goal is resolved, such that when next encountering the same situation, the new production fires without the need to enter a new sub-goal. This type of information can include which operator to select, or how to implement an operator. These rules tend to be atomic, and in nearly all cases can be seen as immediately fully learned. This learning mechanism (chunking) can implement a wide range of learning effects, including long-term declarative memory learning—for long-term declarative information is represented solely as the result of procedural memory.

ACT-R learning involves both declarative and procedural memory. When rules fire they become stronger, and as declarative memories are used more they are strengthened as well. Each production also has an expected gain value based on its probability of success and its cost and the current goal's value. The expected gain is used for conflict resolution; the production with the highest expected gain is selected when several productions are possible matches. The more often the production meets with later success (e.g., the sub-goal ends up being solved), the higher this probability for the rule will become. This strength also influences the activation of the declarative memory items that are matched by the condition of the production, and also the rule execution time.

Each item in declarative memory has an associated activation that changes based upon how often it has been used, and how strongly it is associated with other items that are being used. The more often an item is used, the higher its base level activation will become. The more strongly associated an item is with ones that are being used, the more chance that item has for having its activation raised.

A rule learning mechanism is less often used in ACT-R models, and when it has been used, the resulting rules are typically created in a nascent state such that they have to be created several times before they are fully learned.

B.4 Bibliography for Soar and ACT-R

ai.eecs.umich.edu/soar/, the Soar Group's homepage

act.psy.cmu.edu/, the ACT-R Group's homepage

acs.ist.psu.edu/soar-faq, Soar Frequently Asked Questions list

acs.ist.psu.edu/act-r-faq, ACT-R Frequently Asked Questions list

Jones, G. (1996). The architectures of Soar and ACT-R, and how they model human behaviour. *Artificial Intelligence and Simulation of Behaviour Quarterly*, 96 (Winter), 41-44.

Johnson, T. R. (1997). Control in ACT-R and Soar. In M. Shafto & P. Langley (Eds.), *Proceedings of the Nineteenth Annual Conference of the Cognitive Science Society* (pp. 343-348). Hillsdale, NJ: Erlbaum.

Ritter, F. E. (2002). Soar. In *Encyclopedia of cognitive science*. London: Macmillan.

Glossary of Acronyms and Abbreviations

ABC	A* search with Bounded Costs
ACT-R	Adaptive Control of Thought - Rational
ACT-R/PM	A perceptual-motor component added to ACT-R
AI	Artificial Intelligence
AMBR	Agent-Based Modeling and Behavior Representation project
APEX	A tool for applied human performance modeling developed at NASA
API	Application Programming Interface
ATAL workshops	Architectures, Theories, And Languages Workshop series
BDI architectures	Architectures based on representing Beliefs, Desires, and Intentions
CES	Cognitive Environment Simulation
CHIRP	Confidential Human Factors Incident Reporting Program
CHREST	Chunk Hierarchy and REtrieval STructures
CMAC	Cerebellar Model Arithmetic Computer
CoCoM	Contextual Control Model
COSIMO	COgnitive SIMulation MOdel
CREAM	Cognitive Reliability and Error Analysis Method
DERA	Defence Evaluation and Research Agency (UK)
DCOM	Distributed COmponent Model
DIS	Distributed Interactive Simulation (system)
EPAM	Elementary Perceiver and Memoriser
EPIC	A cognitive architecture based on a production rule interpreter that assumes no cognitive limitations on processing and a set of perceptual motor processors that provide a limitation on cognition.
FLAME	Fuzzy Logic Adaptive Model of Emotions
GAs	Genetic Algorithms

HCI	Human-Computer Interaction
HLA	Higher-Level Architecture
IDM	Individual Data Modeling, modeling based on fitting the behavior of individuals and then aggregating the results, as compared with fitting data aggregated across subjects.
IMPS	Internet-based Multi-agent Problem Solving
JACK	JAVA Agent Compiler and Kernel
JAVA	A procedural language used to support web applications
JFC	JAVA Foundation Classes
JNDI	JAVA Naming and Directory Interface
KBS	Knowledge-Based Systems
LTM	Long-Term Memory
MLP	Multi-Layer Perceptron
ModSAF	Modular Semi-Automated Forces
NDM	Naturalistic Decision Making
ONR	Office of Naval Research
RDM	Rapid Decision Making
RMI	Remote Method Invocation
SDM	Sparse Distributed Memory
SEs	Synthetic Environments
SMOC	Simplified Model Of Cognition
SRG	System Response Generator
STM	Short-Term Memory
UTC	Unified Theory of Cognition

References

- Aasman, J. (1995). *Modelling driver behaviour in Soar*. Leidschendam, The Netherlands: KPN Research.
- Aasman, J., & Michon, J. A. (1992). Multitasking in driving. In J. A. Michon & A. Akyürek (Eds.), *Soar: A cognitive architecture in perspective*. Dordrecht, The Netherlands: Kluwer.
- Agre, P., & Chapman, D. (1987). Pengi: An Implementation of a Theory of Activity. In *Proceedings of the Sixth National Conference on Artificial Intelligence (AAAI-87)* (pp. 268-272). Seattle, WA: Morgan-Kaufman.
- Agre, P. E., & Shragar, J. (1990). Routine evolution as the microgenetic basis of skill acquisition. In *Proceedings of the 12th Annual Conference of the Cognitive Science Society* (pp. 694-701). Hillsdale, NJ: Erlbaum.
- Albus, J. S. (1971). A theory of cerebella function. *Mathematical BioScience*, 10, 25-61.
- Alechina, N., & Logan, B. (2001). State space search with prioritized soft constraints. *Applied Intelligence*, 14 (3), 263-272.
- Amalberti, R., & Deblon, F. (1992). Cognitive modelling of fighter aircraft's control process: A step towards intelligent onboard assistance system. *International Journal of Man-Machine Studies*, 36, 639-671.
- Anderson, J. R. (1983). *The architecture of cognition*. Cambridge, MA: Harvard University Press.
- Anderson, J. R. (1990). *The adaptive character of thought*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R. (1993). *Rules of the mind*. Hillsdale, NJ: Erlbaum.
- Anderson, J. R., Corbett, A. T., Koedinger, K. R., & Pelletier, R. (1995). Cognitive tutors: Lessons learned. *Journal of the Learning Sciences*, 4(2), 167-207.
- Anderson, J. R., Farrell, R., & Sauers, R. (1984). Learning to program in LISP. *Cognitive Science*, 8, 87-129.
- Anderson, J. R., & Lebiere, C. (1998). *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- Anderson, J. R., Matessa, M., & Lebiere, C. (1998). ACT-R: A theory of higher level cognition and its relation to visual attention. *Human-Computer Interaction*, 12, 439-462.
- Angros, R. (1998). Learning what to instruct. In *Proceedings of Soar Workshop 18* (pp. 58-64). Vienna, VA: Explore Reasoning Systems.
- Anzai, Y., & Simon, H. A. (1979). The theory of learning by doing. *Psychological Review*, 86, 124-140.
- Assanie, M., & Laird, J. (1999). Learning by instruction using a constrained natural language interface. In *Proceedings of Soar Workshop 19* (pp. 144-149). Ann Arbor, MI: University of Michigan.

- Baddeley, A. D. (1986). *Working memory*. Oxford, UK: Oxford University Press.
- Baddeley, A. D. (1997). *Human memory: Theory and practice*. Hove, UK: Psychology Press.
- Bartl, C., & Dörner, D. (1998). PSI: A theory of the integration of cognition, emotion and motivation. In F. E. Ritter & R. M. Young (Eds.), *Proceedings of the 2nd European Conference on Cognitive Modelling* (pp. 66-73). Thrumpton, Nottingham, UK: Nottingham University Press.
- Bass, E. J., Baxter, G. D., & Ritter, F. E. (1995). Creating models to control simulations: A generic approach. *AISB Quarterly*, 93, 18-25.
- Baxter, G. D. (1997). *From soup to nuts: Developing a Soar cognitive model of the electronic warfare officer's task* (Working paper No. WP/R3BAIA005/014). Nottingham, UK: University of Nottingham, Department of Psychology, Cognitive Modelling Unit.
- Baxter, G. D., & Ritter, F. E. (1996). *Designing abstract visual perceptual and motor action capabilities for use by cognitive models* (Tech. Rep. No. 36). Nottingham, UK: University of Nottingham, Department of Psychology, ESRC Centre for Research in Development, Instruction and Training.
- Baylor, G. W., & Simon, H. A. (1966). A chess mating combinations program. In *Proceedings of the 1966 Spring Joint Computer Conference*, 28, (pp. 431-447). Boston: AFIPS. Reprinted in H. A. Simon (Ed.). (1979) *Models of thought*. New Haven, CT: Yale University Press.
- Belavkin, R. (2001). The role of emotions in problem solving. In *AISB'01 Symposium on Emotion, Cognition and Affective Computing* (pp. 49-57). Falmer, UK: The Society for the Study of AI and Simulation of Behaviour.
- Belavkin, R., & Ritter, F. E. (2000). Adding a theory of motivation to ACT-R. In *Proceedings of the Seventh Annual ACT-R Workshop* (pp. 133-139). Pittsburgh, PA: Carnegie-Mellon University, Department of Psychology.
- Belavkin, R. V., Ritter, F. E., & Elliman, D. G. (1999). Towards including simple emotions in a cognitive architecture in order to fit children's behaviour better. In *Proceedings of the 1999 Conference of the Cognitive Science Society* (p. 784). Mahwah, NJ: Erlbaum.
- Bigus, J. P., & Bigus, J. (1997). *Constructing intelligent agents in JAVA*. New York: Wiley.
- Bloedorn, G. W., & Downes-Martin, S. G. (1985). *Tank tactics grandmaster* (Rep. No. 5938). Cambridge, MA: BBN Laboratories.
- Boff, K. R., Kaufman, L., & Thomas, J. P. (Eds.). (1986). *Handbook of perception and human performance*. New York: Wiley.
- Boff, K. R., & Lincoln, J. E. (Eds.). (1988). *Engineering data compendium*. Wright-Patterson Air Force Base, OH: Armstrong Aerospace Medical Research Laboratory.

- Brazier, F., Dunin-Keplicz, B., Treur, J., & Verbrugge, R. (1999). Modelling internal dynamic behaviour of BDI agents. In A. Cesto & P. Y. Schobbles (Eds.), *Proceedings of the Third International Workshop on Formal Methods of Agents, MODELAGE '97*, 21. Lecture notes in AI. Berlin: Springer Verlag.
- Brooks, R. (1992). Intelligence without representation. In D. Kirsh (Ed.), *Foundations of artificial intelligence*. Cambridge, MA: MIT Press.
- Burke, E. K., Elliman, D. G., & Weare, R. F. (1995). A hybrid genetic algorithm for highly constrained timetabling problems. In *Proceedings of the 6th International Conference on Genetic Algorithms*. (pp. 605-610). Seattle: Morgan Kaufman.
- Busetta, P., Howden, N., Rönquist, R., & Hodgson, A. (1999a). Structuring BDI agents in functional clusters. In *Proceedings of the Sixth International Workshop on Agent Technologies, Architectures and Languages* (pp. 149-161).
- Busetta, P., Rönquist, R., Hodgson, A., & Lucas, A. (1999b, Jan). JACK intelligent agents —Components for intelligent agents in JAVA. *AgentLink News Letter*, 2. www.agent-software.com.
- Butler, E. (2000). *A dynamical study of the generalised delta rule*. Unpublished doctoral dissertation, University of Nottingham, Nottingham, UK.
- Byrne, M. D. (2001). ACT-R/PM and menu selection: Applying a cognitive architecture to HCI. *International Journal of Human-Computer Studies*, 55, 41-84.
- Byrne, M. D. (1997). *ACT-R Perceptual-Motor (ACT-R/PM) version 1.0b1: A users manual*. Pittsburgh, PA: Carnegie-Mellon University, Department of Psychology. Available through act.psy.cmu.edu.
- Byrne, M. D., & Anderson, J. R. (1998). Perception and action. In J. R. Anderson & C. Lebiere (Eds.), *The atomic components of thought*. Mahwah, NJ: Erlbaum.
- Byrne, M. D., & Bovair, S. (1997). A working memory model of a common procedural error. *Cognitive Science*, 21(1), 31-61.
- Byrne, M. D., Chong, R., Freed, M., Ritter, F. E., & Gray, W. (1999). Symposium on Integrated models of perception, cognition, and action. In *Proceedings of the 1999 Conference of the Cognitive Science Society*, 1. Mahwah, NJ: Erlbaum.
- Cacciabue, P. C., Decortis, F., Drozdowicz, B., Masson, M., & Nordvik, J. (1992). COSIMO: A cognitive simulation model of human decision making and behavior in accident management of complex plants. *IEEE Transactions on Systems, Man, and Cybernetics*, 22(5), 1058-1074.
- Campbell, C., Hull, R., Root, E., & Jackson, L. (1995). Route Planning in CCTT. In *Proceedings of the Fifth Conference on Computer Generated Forces and Behavioral Representation* (pp. 233-244). Orlando, FL: University of Central Florida, Institute for Simulation and Training.
- Ceranowicz, A.Z. (1994, May 13). *ModSAF and Command and Control, Modular Semi-Automated Forces: Recent and Historical Publications*. Cambridge, MA: Loral Advanced Distributed Simulation.

- Ceranowicz, A. (1998). STOW 97 - 99. In *Proceedings of the 7th Conference on Computer Generated Forces and Behavioral Representation* (pp. 3-17). Orlando, FL: University of Central Florida, Division of Continuing Education.
- Chandrasekaran, B., & Josephson, J. R. (1999). Cognitive modeling for simulation goals: A research strategy for computer-generated forces. In *Proceedings of the 8th Computer Generated Forces and Behavioral Representation Conference* (pp. 239-250). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/8th/view-papers.htm.
- Charness, N. (1991). Expertise in chess: The balance between knowledge and search. In K. A. Ericsson & J. Smith (Eds.), *Studies of expertise: Prospects and limits*. Cambridge, UK: Cambridge University Press.
- Charness, N. (1992). The impact of chess research on cognitive science. *Psychological Research*, 54, 4-9.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55-81.
- Chipman, S., & Meyrowitz, A. L. (Eds.). (1993). *Foundations of knowledge acquisition: Cognitive models of complex learning*. Boston, MA: Kluwer.
- Chong, R. (1999). Towards a model of fear in Soar. In *Proceedings of Soar Workshop 19* (pp. 6-9). Ann Arbor, MI: University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop19/talks/proceedings.html.
- Chong, R. S. (2001). Low-level behavioral modeling and the HLA: An EPIC-Soar model of an enroute air-traffic control task. In *Proceedings of the 10th Computer Generated Forces and Behavioral Representation Conference* (pp. 27-35). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/10th
- Chong, R. S., & Laird, J. E. (1997). Identifying dual-task executive process knowledge using EPIC-Soar. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society* (pp. 107-112). Mahwah, NJ: Erlbaum.
- Computer Science and Telecommunications Board, N. R. C. (1997). *Modeling and simulation: Linking entertainment and defense*. Washington, DC: National Academy Press.
- Cooper, R., & Fox, J. (1998). COGENT: A visual design environment for cognitive modelling. *Behavior Research Methods, Instruments and Computers*, 30, 553-564.
- Cottam, H., & Shadbolt, N. R. (1998). Knowledge acquisition for search and rescue planning. *International Journal of Human-Computer Studies*, 48, 449-473.
- Crow, L., & Shadbolt, N. R. (1998). IMPS - Internet agents for knowledge engineering. In B. R. Gaines & M. Musen (Eds.), *Knowledge Acquisition Workshop KAW'98* (pp. 1-19). Banff, Alberta, Canada: SRAG Publications.
- Crowder, R. G. (1976). *Principles of learning and memory*. Hillsdale, NJ: Erlbaum.
- Daily, L. Z., Lovett, M. C., & Reder, L. M. (2001). Modeling individual differences in working memory performance: A source activation account. *Cognitive Science*, 25(3), 315-355.

- Damasio, A. R. (1994). *Descartes' error: Emotion, reason, and the human brain*. New York: Gosset/Putnam Press.
- Daneman, M., & Carpenter, P. A. (1980). Individual differences in working memory and reading. *Journal of Verbal Learning and Verbal Behavior*, 19, 450-466.
- Dawes, R. M. (1988). *Rational choice in an uncertain world*. Orlando, FL: Harcourt Brace Jovanovich.
- Dawes, R. M. (1994). *House of cards: Psychology and psychotherapy built on myth*. New York: Free Press.
- de Groot, A. D. (1946). *Het denken van den schaker (Thought and choice in chess)*. Amsterdam: Noord Hollandsche.
- de Groot, A. D. (1978). *Thought and choice in chess*. The Hague, The Netherlands: Mouton Publishers.
- de Groot, A. D., & Gobet, F. (1996). *Perception and memory in chess*. Assen, The Netherlands: Van Gorcum.
- de Keyser, V., & Woods, D. D. (1990). Fixation errors: Failures to revise situation assessment in dynamic and risky systems. In A. G. Colombo & A. S. de Bustamante (Eds.), *Systems reliability assessment* (pp. 231-251). Dordrecht, The Netherlands: Kluwer.
- Delaney, P. F., Reder, L. M., Staszewski, J. J., & Ritter, F. E. (1998). The strategy specific nature of improvement: The power law applies by strategy within task. *Psychological Science*, 9(1), 1-8.
- Detje, F. (2000). Comparison of the PSI-theory with human behaviour in a complex task. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the Third International Conference on Cognitive Modelling* (pp. 86-93). Veenendaal, The Netherlands: Universal Press.
- Dörner, D. (2000). The simulation of extreme forms of behavior. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the Third International Conference on Cognitive Modelling* (pp. 94-99). Veenendaal, The Netherlands: Universal Press.
- Elkind, J. I., Card, S. K., Hochberg, J., & Huey, B. M. (Eds.). (1989). *Human performance models for computer-aided engineering*. Washington, DC: National Academy Press.
- Elliman, D. G. (1989). Machine recognition of engineering drawings. In G. Bryan (Ed.), *Exploitable UK research for manufacturing industry*. London: Peregrinus Press.
- Elman, J. L. (1991). Distributed representations, simple recurrent networks, and grammatical structure. *Machine Learning*, 7, 194-220.
- Ericsson, K. A., & Kintsch, W. (1995). Long-term working memory. *Psychological Review*, 102, 211-245.
- Ericsson, K. A., & Simon, H. A. (1993). *Protocol analysis: Verbal reports as data*. Cambridge, MA: MIT Press.
- Etzioni, O., & Weld, D. S. (1995). Intelligent agents on the Internet - Fact, fiction, and forecast. *IEEE Expert - Intelligent Systems & Their Applications*, 10(4), 44-49.

- Feigenbaum, E. A., & Simon, H. A. (1962). A theory of the serial position effect. *British Journal of Psychology*, 53, 307-320.
- Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, 8, 305-336.
- Franceschini, R. W., McBride, D. K., & Sheldon, E. (2001). Recreating the Vincennes Incident using affective computer generated forces. In *Proceedings of the 10th Computer Generated Forces and Behavioral Representation Conference*. 10TH-CG-044.doc. Orlando, FL: Division of Continuing Education, University of Central Florida. www.sisostds.org/cgf-br/10th/.
- Franklin, S., & Graesser, A. (1997). Is it an agent, or just a program? A taxonomy for autonomous agents. In J. P. Müller, M. J. Wooldridge, & N. R. Jennings (Eds.), *Intelligent Agents III—Proceedings of the Third International Workshop on Agent Theories, Architectures, and Languages (ATAL-96), Lecture Notes in Artificial Intelligence* (pp. 21-36). Berlin: Springer-Verlag.
- Freed, M., & Remington, R. (2000). Making human-machine system simulation a practical engineering tool: An APEX overview. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the 3rd International Conference on Cognitive Modelling* (pp. 110-117). Veenendaal, The Netherlands: Universal Press.
- Freed, M. A., Shafto, M. G., & Remington, R. W. (1998). Employing simulation to evaluate designs: The APEX approach. In *Proceedings of the 2nd Workshop on Human Error, Safety, and System Development* (pp. 120-132).
- Frese, M., & Altmann, A. (1989). The treatment of errors in training and learning. In L. Bainbridge & S. A. R. Quintanilla (Eds.), *Developing skills with information technology* (pp. 65-86). Chichester, UK: Wiley.
- Geddes, N. D. (1989). *Understanding human intentions through action interpretation*. Unpublished doctoral dissertation, Georgia Institute of Technology, Atlanta.
- Gigley, H. M., & Chipman, S. F. (1999). Productive interdisciplinarity: The challenge that human learning poses to machine learning. In *Proceedings of the 21st Conference of the Cognitive Science Society 2*. Mahwah, NJ: Erlbaum.
- Glover, F., & Laguna, M. (1998). *Tabu search*. Dordrecht, The Netherlands: Kluwer.
- Gluck, K. A., & Pew, R. W. (2001a). Lessons learned and future directions for the AMBR model comparison project. In *Proceedings of the 10th Computer Generated Forces and Behavioral Representation Conference* (pp. 113-121). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/10th/.
- Gluck, K. A., & Pew, R. W. (2001b). Overview of the agent-based modeling and behavior representation (AMBR) model comparison project. In *Proceedings of the 10th Computer Generated Forces and Behavioral Representation Conference* (pp. 3-6). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/10th/.

- Gobet, F. (1997). A pattern-recognition theory of search in expert problem solving. *Thinking and Reasoning*, 3, 291-313.
- Gobet, F. (1998). Expert memory: A comparison of four theories. *Cognition*, 66, 115-152.
- Gobet, F., & Jansen, P. (1994). Towards a chess program based on a model of human memory. In H. J. van den Herik, I. S. Herschberg, & J. E. Uiterwijk (Eds.), *Advances in computer chess 7*. Maastricht, The Netherlands: University of Limburg Press.
- Gobet, F., Richman, H., Staszewski, J., & Simon, H. A. (1997). Goals, representations, and strategies in a concept attainment task: The EPAM model. *The Psychology of Learning and Motivation*, 37, 265-290.
- Gobet, F., & Ritter, F. E. (2000). Individual Data Analysis and Unified Theories of Cognition: A methodological proposal. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the 3rd International Conference on Cognitive Modelling* (pp. 150-157). Veenendaal, The Netherlands: Universal Press.
- Gobet, F., & Simon, H. A. (1996a). The roles of recognition processes and look-ahead search in time-constrained expert problem solving: Evidence from grandmaster level chess. *Psychological Science*, 7, 52-55.
- Gobet, F., & Simon, H. A. (1996b). Templates in chess memory: A mechanism for recalling several boards. *Cognitive Psychology*, 31, 1-40.
- Gobet, F., & Simon, H. A. (2000). Five seconds or sixty? Presentation time in expert memory. *Cognitive Science*, 24, 651-682.
- Gobet, F., & Simon, H. A. (2001). Human learning in game playing. In M. Kubat, M. Fürnkranz, & J. Fürnkranz (Eds.), *Machines that learn to play games. NOVA science, Advances in Computation: Theory and Practice* (Vol. 8, pp. 61-80). Huntington, NY: Nova Science.
- Goldberg, D. E. (1989). *Genetic algorithms in search, optimization, and machine learning*. Reading, MA: Addison-Wesley.
- Gratch, J. (1998). Planning in Soar. In *Proceedings of Soar Workshop 18* (pp. 17-25). Vienna, VA: Explore Reasoning Systems.
- Gratch, J. (1999). Why you should buy an emotional planner. In *Proceedings of Soar Workshop 19* (pp. 1-3). University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop19/talks/proceedings.html
- Gratch, J., & Marsella, S. (2001). Modeling emotions in the Mission Rehearsal Exercise. In *Proceedings of the 10th Computer Generated Forces and Behavioural Representation Conference* (10TH—CG-057). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/10th/.
- Green, R. (1990). Human error on the flight deck. *Philosophical Transactions of the Royal Society London, Series B*, 327, 503-512.
- Greenberg, A. D., Small, R. L., Zenyah, J. P., & Skidmore, M. D. (1995). Monitoring for hazard in flight management systems. *European Journal of Operations Research*, 84, 5-24.

- Grimes, C., Picton, P. D., & Elliman, D. G. (1996). A neural network position independent multiple pattern recogniser. *Artificial Intelligence in Engineering*, 10, 117-126.
- Hart, P. E., Nilsson, N. J., & Raphael, B. (1968). A formal basis for the heuristic determination of minimum cost paths. *IEEE Transactions on Systems Science and Cybernetics*, SSC-4(2), 100-107.
- Heise, D. R. (1989). Modeling event structures. *Journal of Mathematical Sociology*, 14(2-3), 139-169.
- Heise, D. R., & Lewis, E. (1991). *Introduction to ETHNO*. Dubuque, IA: William. C. Brown.
- Hill, R. (1999). Modeling perceptual attention in virtual humans. In *Proceedings of Soar Workshop 19*. 99-102. Ann Arbor, MI: University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop19/talks/proceedings.html
- Hoffman, R. R. (1987). The problem of extracting the knowledge of experts from the perspective of experimental psychology. *AI Magazine*, 8, 53-67.
- Hoffman, R. R. (1992). *The psychology of expertise. Cognitive research and empirical AI*. New York: Springer-Verlag.
- Hoffman, R. R., Crandall, B., & Shadbolt, N. R. (1998). Use of the Critical Decision Method to elicit expert knowledge: A case study in the methodology of Cognitive Task Analysis. *Human Factors*, 40, 254-276.
- Hoffman, R. R., & Shadbolt, N. R. (1995). *A review of "Naturalistic Decision-Making" research on the critical decision method of knowledge elicitation and the recognition priming model of decision-making* (Report commissioned by the Defence Research Agency, Project No ASF/2188/U). Nottingham, UK: University of Nottingham, School of Psychology, AI Group.
- Hoffman, R. R., & Shadbolt, N. R. (1996). *Facilitating the acquisition of expertise in domains involving perceptual skill, mental workload, and situation awareness* (Report commissioned by the Defence Research Agency, Project No ASF/2819U).
- Holland, J. H., Holyoak, K., Nisbett, R., & Thagard, P. (1986). *Induction: Processes of inference, learning, and discovery*. Cambridge, MA: MIT Press.
- Hollands, R., Denby, B., & Brooks, G. (1999). SAFE-VR—A virtual reality safety training system. In *Proceedings of 28th International Symposium on Computer Applications in the Minerals Industry* (pp. 787-794). Golden, CO: Colorado School of Mines. http://www.mines.edu/outreach.cont_ed/apcom.htm
- Hollnagel, E. (1993). *Human reliability analysis: Context and control*. London: Academic.
- Hollnagel, E. (1998). *Cognitive reliability and error assessment method*. Oxford, UK: Elsevier.
- Hollnagel, E., & Cacciabue, C. (1991, Sep 2-6). Cognitive modelling in system simulation. In *Proceedings of the Third European Conference on Cognitive Science Approaches to Process Control*, 1-29, Cardiff, UK.

- Howes, A. (1993). Recognition-based problem solving. In *Proceedings of the Fifteenth Annual Meeting of the Cognitive Science Society* (pp. 551-556). Hillsdale, NJ: Erlbaum.
- Howes, A., & Young, R. M. (1996). Learning consistent, interactive, and meaningful task-action mappings: A computational model. *Cognitive Science*, 20(3), 301-356.
- Hudlicka, E. (1997). *Modeling behavior moderators in military human performance models* (Tech. Rep. No. 9716). Lincoln, MA: Psychometrix. (Substantial portions of this report appear in Pew, R. S., & Mavor, A. S. [Eds.]. [1998]. *Representing Human Behavior in Military Simulations* [NRC Panel Report]. Washington, DC: National Academy Press.)
- Hudlicka, E., & Fellous, J.-M. (1996). Review of computational models of emotion (Tech. Rep. No. 9612). Arlington, MA: Psychometrix.
- Huffman, S. B., & Laird, J. E. (1995). Flexibly instructable agents. *Journal of AI Research*, 3, 271-324.
- Jacobs, N., & Shea, R. (1996). *The role of JAVA in InfoSleuth: Agent-based exploitation of heterogeneous information resources* (Tech. Rep. No. MCC-INSL-018-96). Presented at the IntraNet96 JAVA Developers Conference. Formerly available online at www.mcc.com/projects/infosleuth/publications/intranet-java.html.
- Jansen, P. J. (1992). *Using knowledge about the opponent in game-tree search*. Unpublished doctoral dissertation (CMU-CS-92-192), Carnegie-Mellon University, Pittsburgh.
- John, B. E., & Kieras, D. E. (1996). Using GOMS for user interface design and evaluation: Which technique? *ACM Transactions on Computer-Human Interaction*, 3(4), 287-319.
- John, B. E., Vera, A. H., & Newell, A. (1994). Towards real-time GOMS: A model of expert behavior in a highly interactive task. *Behavior and Information Technology*, 13, 255-267.
- John, B., Vera, A., Matessa, M., Freed, M., & Remington, R. (2002). Automating CPM-GOMS. In *Proceedings of the CHI'02 Conference on Human Factors in Computer Systems*. New York: ACM.
- Jones, B. (2001). Soar General Input/Output Interface (SGIO). In *Proceedings of 21st Soar Workshop* (pp. 26-28). Ann Arbor, MI: The University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop21/talks/author-index.html
- Jones, B. L., Bauman, J. J., & Laird, J. E. (2001). Visual Soar. In *Proceedings of 21st Soar Workshop* (pp. 71-72). Ann Arbor, MI: The University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop21/talks/author-index.html
- Jones, G. (1999a). *Testing mechanisms of development within a computational framework*. Unpublished doctoral dissertation, University of Nottingham.
- Jones, G., & Ritter, F. E. (1998). Initial explorations of simulating cognitive and perceptual development by modifying architectures. In *Proceedings of the 20th Annual Conference of the Cognitive Science Society* (pp. 543-548). Mahwah, NJ: Erlbaum.
- Jones, G., Ritter, F. E., & Wood, D. J. (2000). Using a cognitive architecture to examine what develops. *Psychological Science*, 11(2), 93-100.

- Jones, R. (1998). Modeling pilot fatigue with a synthetic behavior model. In *Proceedings of the 7th Conference on Computer Generated Forces and Behavioral Representation* (pp. 349-357). Orlando, FL: University of Central Florida, Division of Continuing Education.
- Jones, R. M. (1999b). Graphic visualization of situation awareness and mental state for intelligent computer-generated forces. In *Proceedings of the Eighth Conference on Computer Generated Forces and Behavioral Representations* (pp. 219-222). Orlando, FL: University of Central Florida, Division of Continuing Education.
- Jones, R. M., Laird, J. E., Nielsen, P. E., Coulter, K. J., Kenny, P., & Koss, F. V. (1999). Automated intelligent pilots for combat flight simulation. *AI Magazine*, 20(1), 27-41.
- Just, M. A., & Carpenter, P. A. (1992). A capacity theory of comprehension: Individual differences in working memory. *Psychological Review*, 99, 122-149.
- Kalus, T., & Hirst, T. (1999). ViSoar. In *Proceedings of Soar Workshop 19* (pp. 70-73). Ann Arbor, MI: The University of Michigan Soar Group. www.dcs.port.ac.uk/~hirsta/visoarx.htm
- Kanerva, P. (1988). *Sparse distributed memory*. Cambridge, MA: MIT Press.
- Kieras, D. (1985). The role of cognitive simulation models in the development of advanced training and testing systems. In N. Frederiksen, R. Glaser, A. Lesgold, & M. G. Shafto (Eds.), *Diagnostic monitoring of skill and knowledge acquisition* (pp. 365-394). Hillsdale, NJ: Erlbaum.
- Kieras, D. E., & Meyer, D. E. (1997). An overview of the EPIC architecture for cognition and performance with application to human-computer interaction. *Human-Computer Interaction*, 12, 391-438.
- Kintsch, W. (1998). *Comprehension: A paradigm for cognition*. Cambridge, UK: Cambridge University Press.
- Kitajima, M., & Polson, P. G. (1996). A comprehension-based model of exploration. In M. J. Tauber, V. Bellotti, R. Jeffries, J. D. MacKinlay, & J. Nielsen (Eds.), *Proceedings of the CHI '96 Conference on Human Factors in Computer Systems* (pp. 324-331). New York: ACM.
- Kitajima, M., Soto, R., & Polson, P. G. (1998). LICAI+: A comprehension-based model of the recall of action sequences. In F. E. Ritter & R. M. Young (Eds.), *Proceedings of the 2nd European Conference on Cognitive Modelling* (pp. 82-89). Thrumpton, Nottingham, UK: Nottingham University Press.
- Klein, G. A. (1997). *Recognition-primed decision making: Looking back, looking forward*. Hillsdale, NJ: Erlbaum.
- Koriat, A., & Lieblich, I. (1974). What does a person in a "TOT" state know that a person in a "don't know" state doesn't know? *Memory & Cognition*, 2(4), 647-655.
- Kosslyn, S. M., & Koenig, O. (1992). *Wet mind*. New York: Free Press.

- Kuk, G., Arnold, M., & Ritter, F. E. (1999). Using event history analysis to model the impact of workload on an air traffic tactical controller's operations. *Ergonomics*, 42(9), 1133-1148.
- Laird, J. E. (1999). Visual Soar. In *Proceedings of Soar Workshop 19* (pp. 99-102). Ann Arbor, MI: University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop19/talks/proceedings.html
- Laird, J. E., Coulter, K., Jones, R., Kenny, P., Koss, F., & Nielsen, P. (1997). *Review of Soar/FWA participation in STOW-97*. ai.eecs.umich.edu/ifor/stow-review.html
- Laird, J. E., & Duchi, J. C. (2000). Creating human-like synthetic characters with multiple skill levels: A case study using the Soar Quakebot. In *Simulating Human Agents, Papers from the 2000 AAAI Fall Symposium* (pp. 75-79). Menlo Park, CA: AAAI Press.
- Laird, J. E., Newell, A., & Rosenbloom, P. S. (1987). Soar: An architecture for general intelligence. *Artificial Intelligence*, 33(1), 1-64.
- Laird, J. E., Rosenbloom, P. S., & Newell, A. (1986). Chunking in Soar: The anatomy of a general learning mechanism. *Machine Learning*, 1(1), 11-46.
- Lane, P. C. R., Cheng, P. C.-H., & Gobet, F. (1999). Problem solving with diagrams: Modelling the implicit learning of perceptual information (Tech. Rep. No. 59). University of Nottingham, UK: Department of Psychology, ESRC Centre for Research in Development, Instruction and Training.
- Langley, P., & Ohlsson, S. (1984). Automated cognitive modeling. In *Proceedings of the 2nd National Conference on Artificial Intelligence (AAAI-84)* (pp. 193-197). Los Altos, CA: Morgan Kaufman.
- Larkin, J. H. (1981). Enriching formal knowledge: A model for learning to solve textbook physics problems. In J. R. Anderson (Ed.), *Cognitive skills and their acquisition* (pp. 311-334). Hillsdale, NJ: Erlbaum.
- Larkin, J. H., McDermott, J., Simon, D. P., & Simon, H. A. (1980). Expert and novice performance in solving physics problems. *Science*, 208, 1335-1342.
- Laughery, K. R., & Corker, K. M. (1997). Computer modeling and simulation of human/system performance. In G. Salvendy (Ed.), *Handbook of human factors*. New York: Wiley.
- Lebiere, C. (2001). Multi-tasking and cognitive workload in an ACT-R model of a simplified air traffic control task. In *Proceedings of the 10th Computer Generated Forces and Behavioral Representation Conference* (pp. 91-98). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/10th/.
- Lebiere, C., Anderson, J. R., & Reder, L. M. (1994). Error modeling in the ACT-R production system. In *Proceedings of the 16th Annual Conference of the Cognitive Science Society* (pp. 555-559). Hillsdale, NJ: Erlbaum.

- Lesgold, A. M., Lajoie, S., Bunzon, M., & Eggan, E. (1992). SHERLOCK: A coached practice environment for an electronics troubleshooting job. In J. Larkin, R. Chabay, & C. Scheftic (Eds.), *Computer assisted instruction and intelligent tutoring systems: Establishing communication and collaboration*. Hillsdale, NJ: Erlbaum.
- Lewis, R. L. (1993). *An architecturally-based theory of human sentence comprehension*. Unpublished doctoral dissertation (CMU-CS-93-226), Carnegie-Mellon University, Pittsburgh.
- Logan, B. (1998). Classifying agent systems. In J. Baxter & B. Logan (Eds.), *Software Tools for Developing Agents: Papers from the 1998 Workshop* (Tech. Rep. WS-98-10 11-21). Menlo Park, CA: AAAI Press. www.cs.nott.ac.uk/~bsl/aaai-98/papers/logan.pdf.
- Logan, B., & Alechina, N. (1998). A* with bounded costs. In *Proceedings of the Fifteenth National Conference on Artificial Intelligence (AAAI-98)* (pp. 444-449). Menlo Park, CA & Cambridge, MA: AAAI Press/MIT Press.
- Logan, B., & Sloman, A. (1998). *Qualitative decision support using prioritised soft constraints* (Tech. Rep. CSRP-98-14). Birmingham, UK: University of Birmingham, School of Computer Science.
- Lonsdale, P. R., & Ritter, F. E. (2000). Soar/Tcl-PM: Extending the Soar architecture to include a widely applicable virtual eye and hand. In N. Taatgen & J. Aasman (Eds.), *Proceedings of the 3rd International Conference on Cognitive Modelling* (pp. 202-209). Veenendaal, The Netherlands: Universal Press.
- Loral (1995, Apr 27). "ModSAF User Manual," Version 1.5.1 (LADS Document Number 95009 v1.1. Cambridge, MA: Loral Advanced Distributed Simulation.
- Lovett, M. C., Daily, L. Z., & Reder, L. M. (2000). A source activation theory of working memory: cross-task prediction of performance in ACT-R. *Journal of Cognitive Systems Research, 1*, 99-118.
- Lucas, A., & Goss, S. (1999, Feb 8-10). The potential for intelligent software agents in defence simulation. In *Proceedings of IDC99, IEEE Symposium on Information, Decision and Control, Adelaide, Australia*, 579-584.
- Mach, E. (1976). *Knowledge and error*. Boston: Reidel Publishing Company. (Translation of the 1905 edition by Thomas J. McCormack and Paul Foulkes. Vienna Circle Collection, Vol. 3. Dordrecht, The Netherlands.)
- McClelland, J. L., & Rumelhart, D. E. (1988). *Explorations in parallel distributed processing: A handbook of models, programs, and exercises*. Cambridge, MA: MIT Press.
- Miller, C. S., & Laird, J. E. (1996). Accounting for graded performance within a discrete search framework. *Cognitive Science, 20*, 499-537.
- Miller, G. A. (1956). The magic number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review, 63*, 81-97.
- Misker, J., Taatgen, N. A., & Aasman, J. (2001). Validating a tool for simulating user interaction. In *Proceedings of ICCM-2001-Fourth International Conference on Cognitive Modeling* (pp. 163-168). Mahwah, NJ: Erlbaum.

- Miyake, A., & Shah, P. (Eds.). (1999). *Models of working memory: Mechanisms of active maintenance and executive control*. New York: Cambridge University Press.
- Moffat, D. C., & Frijda, N. H. (2000). Functional models of emotion. In G. Hatano, N. Okada, & H. Tanabe (Eds.), *Affective minds* (pp. 59-68). Amsterdam: Elsevier.
- Monk, A., Sasse, A., & Crerar, A. (1999). *Affective computing: The role of emotion in human-computer interaction*. British HCI Group one-day meeting in conjunction with University College London (collection of abstracts). www-users.york.ac.uk/~am1/affective.html.
- Neisser, U. (1976). *Cognition and reality*. San Francisco: Freeman.
- Nelson, G., Lehman, J. F., & John, B. E. (1994). Integrating cognitive capabilities in a real-time task. In *Proceedings of 16th Annual Conference of the Cognitive Science Society* (pp. 658-663). Hillsdale, NJ: Erlbaum.
- Nerb, J., Spada, H., & Ernst, A. M. (1997). A cognitive model of agents in a commons dilemma. In *Proceedings of the 19th Annual Conference of the Cognitive Science Society* (pp. 560-565). Mahwah, NJ: Erlbaum.
- Newell, A. (1968). On the analysis of human problem solving protocols. In J. C. Gardin & B. Jaulin (Eds.), *Calcul et formalisation dans les sciences de l'homme* (pp. 145-185). Paris: Centre National de la Recherche Scientifique.
- Newell, A. (1973). You can't play 20 questions with nature and win. In W. G. Chase (Ed.), *Visual information processing* (pp. 283-308). New York, NY: Academic.
- Newell, A. (1982). The knowledge level. *Artificial Intelligence*, 18, 87-127.
- Newell, A. (1990). *Unified theories of cognition*. Cambridge, MA: Harvard University Press.
- Newell, A., Shaw, J. C., & Simon, H. A. (1958). Chess-playing programs and the problem of complexity. *IBM Journal of Research and Development*, 2, 320-335.
- Newell, A., & Simon, H. A. (1972). *Human problem solving*. Englewood Cliffs, NJ: Prentice-Hall.
- Nielsen, T. E., & Kirsner, K. (1994). A challenge for Soar: Modeling proactive expertise in a complex dynamic environment. In *Singapore International Conference on Intelligent Systems (SPICIS-94)*. B79-B84.
- Nwana, H. S. (1996). Software agents: An overview. *The Knowledge Engineering Review*, 11(3), 205-244.
- O'Hara, K., & Shadbolt, N. R. (1998). Generalised directive models: Integrating model development and knowledge acquisition. *International Journal of Human-Computer Studies*, 49, 497-522.
- Ohlsson, S. (1992). Artificial instruction: A method for relating learning theory to instructional design. In M. Jones & P. H. Winne (Eds.), *Adaptive learning environments: Foundations and frontiers* (pp. 55-83). Berlin: Springer-Verlag.

- Ong, R., & Ritter, F. E. (1995). Mechanisms for routinely tying cognitive models to interactive simulations. In *HCI International '95: Poster sessions abridged proceedings*, 84. Osaka, Japan: Musashi Institute of Technology, Department of Industrial Engineering.
- Ong, R. L. (1995). *Mongsu 2.0: Socket utility for hooking up Soar to simulations with sockets*. Available via ritter.ist.psu.edu/nottingham/ccc/
- Payne, S. J. (1991). Display-based action at the user interface. *International Journal of Man-Machine Studies*, 35, 275-289.
- Pearl, J. (1982). A* - An algorithm using search effort estimates. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 4(4), 392-399.
- Pew, R. W., & Mavor, A. S. (Eds.). (1998). *Modeling human and organizational behavior: Application to military simulations*. Washington, DC: National Academy Press. books.nap.edu/catalog/6173.html.
- Piaget, J. (1954). *The construction of reality in the child*. New York: Basic.
- Picard, R. A. (1999). Response to Sloman's review. *The AI Magazine*, 20(1), 134-137.
- Picard, R. W. (1997). *Affective computing*. Cambridge, MA: MIT Press.
- Pitrat, J. (1977). A chess combinations program which uses plans. *Artificial Intelligence*, 8, 275-321.
- Popper, K. R. (1959). *The logic of scientific discovery*. New York: Basic.
- Pylyshyn, Z. (1999). Is vision continuous with cognition? The case for cognitive impenetrability of visual perception [lead article and responses]. *Behavioural and Brain Sciences*, 22(3), 341-365.
- Quinlan, R. (1983). Learning efficient classification procedures and their application to chess end games. In R. S. Michalski, J. G. Carbonell, & T. M. Mitchell (Eds.), *Machine learning: An artificial intelligence approach*. Palo Alto, CA: Tioga.
- Rasmussen, J. (1983). Skills, rules, knowledge: signals, signs and symbols and other distinctions in human performance models. *IEEE Transactions: Systems, Man & Cybernetics, SMC-13*, 257-267.
- Rauterberg, M. (1993). AMME: An automatic mental model evaluation to analyse user behavior traced in a finite, discrete state space. *Ergonomics*, 36(11), 1369-1380.
- Reason, J. (1990). *Human error*. Cambridge, UK: Cambridge University Press.
- Richman, H. B., Gobet, F., Staszewski, J. J., & Simon, H. A. (1996). Perceptual and memory processes in the acquisition of expert performance: The EPAM model. In K. A. Ericsson (Ed.), *The road to excellence*. Mahwah, NJ: Erlbaum.
- Richman, H. B., & Simon, H. A. (1989). Context effects in letter perception: Comparison of two theories. *Psychological Review*, 96, 417-432.
- Richman, H. B., Staszewski, J. J., & Simon, H. A. (1995). Simulation of expert memory with EPAM IV. *Psychological Review*, 102, 305-330.

- Rieman, J., Young, R. M., & Howes, A. (1996). A dual-space model of iteratively deepening exploratory learning. *International Journal of Human-Computer Studies*, 44, 743-775.
- Ritter, F., & Feurzeig, W. (1988). Teaching real-time tactical thinking. In J. Psotka, L. D. Massey, & S. A. Mutter (Eds.), *Intelligent tutoring systems: Lessons learned* (pp. 285-301). Hillsdale, NJ: Erlbaum.
- Ritter, F. E. (1991). Towards fair comparisons of connectionist algorithms through automatically generated parameter sets. In *Proceedings of the 13th Annual Conference of the Cognitive Science Society* (pp. 877-881). Hillsdale, NJ: Erlbaum.
- Ritter, F. E. (1993a). TBPA: A methodology and software environment for testing process models' sequential predictions with protocols (Tech. Rep. No. CMU-CS-93-101). Pittsburgh, PA: Carnegie Mellon University, School of Computer Science.
- Ritter, F. E. (1993b). Three types of emotional effects that will occur in cognitive architectures. In *Workshop on architectures underlying motivation and emotion (WAUME93)*. Birmingham, UK: University of Birmingham, School of Computer Science and Centre for Research in Cognitive Science. acs.ist.psu.edu/papers/ritter93e.pdf
- Ritter, F. E. (2000). A role for cognitive architectures: Guiding user interface design, contribution to the Applications of Cognitive Architectures panel (slides). In *Proceedings of the Seventh Annual ACT-R Workshop* (pp. 85-91). Pittsburgh, PA: Carnegie-Mellon University, Department of Psychology.
- Ritter, F. E., Baxter, G. D., Jones, G., & Young, R. M. (2000). Supporting cognitive models as users. *ACM Transactions on Computer-Human Interaction*, 7(2), 141-173.
- Ritter, F. E., & Bibby, P. (2001). Modeling how and when learning happens in a simple fault-finding task. In *Proceedings of ICCM-2001-Fourth International Conference on Cognitive Modeling* (pp. 187-192). Mahwah, NJ: Erlbaum.
- Ritter, F. E., Jones, G., Baxter, G. D., & Young, R. M. (1998a). Lessons from using models of attention and interaction. In *5th ACT-R Workshop* (pp. 117-123). Pittsburgh, PA: Carnegie Mellon University, Psychology Department.
- Ritter, F. E., Jones, R. M., & Baxter, G. D. (1998b). Reusable models and graphical interfaces: Realising the potential of a unified theory of cognition. In U. Schmid, J. Krems, & F. Wysotzki (Eds.), *Mind modeling-A cognitive science approach to reasoning, learning and discovery* (pp. 83-109). Lengerich, Germany: Pabst Scientific Publishing.
- Ritter, F. E., & Larkin, J. H. (1994). Using process models to summarize sequences of human actions. *Human-Computer Interaction*, 9(3), 345-383.
- Ritter, F. E., & Major, N. P. (1995). Useful mechanisms for developing simulations for cognitive models. *AISB Quarterly*, 91(Spring), 7-18.
- Ritter, F. E., & Wallach, D. P. (1998). Models of two-person games in ACT-R and Soar. In *Proceedings of the Second European Conference on Cognitive Modelling* (pp. 202-203). Nottingham, UK: Nottingham University Press.

- Ritter, F. E., & Young, R. M. (1999). Moving the Psychological Soar Tutorial to HTML: An example of using the Web to assist learning. In D. Peterson, R. J. Stevenson, & R. M. Young (Eds.), *Proceedings of the AISB '99 Workshop on Issues in Teaching Cognitive Science to Undergraduates* (pp. 23-24). Brighton, UK: Society for the Study of Artificial Intelligence and Simulation of Behaviour.
- Ritter, F. E., & Young, R. M. (2001). Embodied models as simulated users: Introduction to this special issue on using cognitive models to improve interface design. *International Journal of Human-Computer Studies*, 55, 1-14.
- Rosenbloom, P. (1998). Emotion in Soar. In *Proceedings of Soar Workshop 18* (pp. 26-28). Vienna, VA: Explore Reasoning Systems.
- Rouse, W. B. (1980). *Systems engineering models of human-machine interaction*. New York: Elsevier.
- Roytam, S. (2001). Soar Lint: Static testing for Soar. In *Proceedings of 21st Soar Workshop* (pp. 73-73). Ann Arbor, MI: The University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop21/talks/author-index.html.
- Saariluoma, P. (1990). Apperception and restructuring in chess player's problem solving. In K. J. Gilhooly, M. T. G. Keane, R. H. Logie, & G. Erdos (Eds.), *Lines of thinking*. New York: Wiley.
- Salvucci, D. (2001). Predicting the effects of in-car interface use on driver performance: An integrated model approach. *International Journal of Human-Computer Studies*, 55, 85-107.
- Sanderson, P. M., & Fisher, C. A. (1994). Exploratory sequential data analysis: Foundations. *Human-Computer Interaction*, 9(3&4), 251-317.
- Sanderson, P. M., McNeese, M. D., & Zaff, B. S. (1994). Handling complex real-world data with two cognitive engineering tools: COGENT and MacSHAPA. *Behavior Research Methods, Instruments, & Computers*, 26(2), 117-124.
- Schofield, D., & Denby, B. (1995). Visualizing the risk—Virtual reality training for LHDs. *Engineering and Mining Journal*, 196(2), 39-40.
- Schraagen, J. M., Chipman, S. F., & Shalin, V. L. (Eds.). (2000). *Cognitive task analysis*. Mahwah, NJ: Erlbaum.
- Schreiber, G., Akkermans, H., Anjewierden, A., de Hoog, R., Shadbolt, N. R., & Wielinga, B. (2000). *Knowledge Engineering and Management*. Cambridge, MA: MIT Press.
- Schunn, C. D., Reder, L. M., Nhouyvanisvong, A., Richards, D. R., & Stroffolino, P. J. (1997). To calculate or not calculate: A source activation confusion (SAC) model of problem-familiarity's role in strategy selection. *Journal of Experimental Psychology: Learning, Memory, & Cognition*, 23, 1-27.
- Schwamb, K. B. (1998). Soar at ERS. In *Proceedings of Soar Workshop 18* (pp. 41-46). Vienna, VA: Explore Reasoning Systems.

- Schwamb, K. B., Koss, F. V., & Keirse, D. (1994). Working with ModSAF: Interfaces for programs and users. In *Proceedings of the Fourth Conference on Computer Generated Forces and Behavioral Representation* (Tech. Rep. IST-TR-94-12). Orlando, FL: University of Central Florida, Institute for Simulation and Training.
- Seif El-Nasr, M., Yen, J., & Ioerger, T. R. (2000). FLAME - Fuzzy Logic Adaptive Model of Emotions. *International Journal of Autonomous Agents and Multi-Agent Systems*, 3(3), 1-39.
- Sekuler, R., & Blake, R. (1994). *Perception* (3rd ed.). New York: Knopf.
- Senders, J. W., & Moray, N. P. (1991). *Human error: Cause, prediction, and reduction*. Hillsdale, NJ: Erlbaum.
- Shadbolt, N. R., & Burton, A. M. (1995). Knowledge elicitation: A systematic approach. In J. R. Wilson & E. N. Corlett (Eds.), *Evaluation of human work: A practical ergonomics methodology* (pp. 406-440). London: Taylor & Francis.
- Shadbolt, N. R., & O'Hara, K. (1997). Model-based expert systems and the explanation of expertise. In P. Feltovich, K. Ford, & R. Hoffman (Eds.), *Expertise in context*. Cambridge, MA: AAAI and MIT Press.
- Siegler, R. S. (1987). The perils of averaging data over strategies: An example from children's addition. *Journal of Experimental Psychology*, 115, 250-264.
- Simon, H. A. (1955). Prediction and hindsight as confirmatory evidence. *Philosophy of Science*, 22, 227-230.
- Simon, H. A. (1974). How big is a chunk? *Science*, 183, 482-488.
- Simon, H. A., & Chase, W. G. (1973). Skill in chess. *American Scientist*, 61, 393-403.
- Singley, M. K., & Anderson, J. R. (1989). *The transfer of cognitive skill*. Cambridge, MA: Harvard University Press.
- Skinner, B. F. (1957). *Verbal behavior*. Englewood Cliffs, NJ: Prentice-Hall.
- Sloman, A. (1998a). Damasio, Descartes, alarms and meta-management. In *Proceedings of the International Conference on Systems, Man, and Cybernetics (SMC98)* (pp. 2652-2657). Piscataway, NJ: IEEE.
- Sloman, A. (1998b). What's an AI toolkit for? In B. Logan & J. Baxter (Eds.), *Proceedings of the AAAI-98 Workshop on Software Tools for Developing Agents* (pp. 1-10). Menlo Park, CA: AAAI Press.
- Sloman, A. (1999). Review of *Affective Computing* by Rosalind Picard (1997). *The AI Magazine*, 20(1), 127-133. (followed by reply by Picard)
- Sloman, A. (2000). Architectural requirements for human-like agents both natural and artificial (What sorts of machines can love?). In K. Dautenhahn (Ed.), *Human cognition and social agent technology. Advances in consciousness research* (pp. 163-195). Amsterdam: John Benjamins.

- Sloman, A., & Logan, B. (1999). Building cognitively rich agents using the Sim_Agent toolkit. *Communications of the Association of Computing Machinery*, 42(3), 71-77.
- St. Amant, R., & Riedl, M. O. (2001). A perception/action substrate for cognitive modeling in HCI. *International Journal of Human-Computer Studies*, 55, 15-39.
- Stern, Ray, R., & Quigley, K. (2001). *Psychophysiological recording*. Oxford, UK: Oxford University Press.
- Sun, R., & Ling, C. X. (1998). Computational cognitive modeling, the source of power, and other related issues. *AI Magazine*, 19(2), 113-120.
- Sun, R., Merrill, E., & Peterson, T. (1998). Skill learning using a bottom-up hybrid model. In F. E. Ritter & R. M. Young (Eds.), *Proceedings of the 2nd European Conference on Cognitive Modelling* (pp. 23-29). Thrumpton, Nottingham, UK: Nottingham University Press.
- Synthetic Environments Management Board (1998). *Synthetic environments and simulation: The report of the Defence and Aerospace Foresight Panel Technology Working Party*. Crossgates, Leeds, UK: Vickers Defence Systems.
- Tabachneck-Schijf, H. J. M., Leonardo, A. M., & Simon, H. A. (1997). CaMeRa: A computational model of multiple representations. *Cognitive Science*, 21(3), 305-350.
- Tambe, M., Johnson, W. L., Jones, R. M., Koss, F., Laird, J. E., Rosenbloom, P. S., & Schwamb, K. (1995). Intelligent agents for interactive simulation environments. *AI Magazine*, 16(1), 15-40.
- Tenney, Y. J., & Spector, S. L. (2001). Comparisons of HBR models with human-in-the-loop performance in a simplified air traffic control simulation with and without HLA protocols: Task simulation, human data and results. In *Proceedings of the 10th Computer Generated Forces and Behavioral Representation Conference* (pp. 15-26). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/10th/view-papers.htm.
- van Lent, M. (1999). Learning by observation in complex domains. In *Proceedings of Soar Workshop 19* (pp. 70-73). Ann Arbor, MI: University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop19/talks/proceedings.html
- van Lent, M., & Laird, J. (1999). Learning by observation in a tactical air combat domain. In *Proceedings of the 8th Computer Generated Forces and Behavioral Representation Conference* (pp. 239-250). Orlando, FL: University of Central Florida, Division of Continuing Education. www.sisostds.org/cgf-br/8th/view-papers.htm
- VanLehn, K., & Garlick, S. (1987). Cirrus: An automated protocol analysis tool. In P. Langley (Ed.), *Fourth Machine Learning Workshop* (pp. 205-217). Los Altos, CA: Morgan Kaufman.
- vanSomeren, M. W., Barnard, Y. F., & Sandberg, J. A. C. (1994). *The Think Aloud Method: A practical guide to modelling cognitive processes*. London: Academic.
- Vicente, K. (1998). *Cognitive work analysis*. Mahwah, NJ: Erlbaum.

- Wallace, C. (2001). Soar API: Designing for speed. In *Proceedings of 21st Soar Workshop* (pp. 16-17). Ann Arbor, MI: The University of Michigan SOAR Group. ai.eecs.umich.edu/soar/workshop21/talks/author-index.html
- Wang, H., Johnson, T. R., & Zhang, J. (1998). UEcho: A model of uncertainty management in human abductive reasoning. In M. A. Gernsbacher & S. J. Derry (Eds.), *Proceedings of the Twentieth Annual Conference of the Cognitive Science Society* (pp. 1113-1129). Mahwah, NJ: Erlbaum.
- Whalen, P. J. (1999). Fear, vigilance, and ambiguity: Initial neuroimaging studies of the human amygdala. *Current Directions*, 7(6), 177-188.
- Wickens, C. D. (1992). *Engineering psychology and human performance* (2nd ed.). New York: Harper Collins.
- Wielinga, B. J., Schreiber, A. T., & Breuker, J. A. (1992). KADS: A modelling approach to knowledge engineering. *Knowledge Acquisition Journal*, 4(1), 5-53.
- Wilkins, D. (1980). Using pattern and plans in chess. *Artificial Intelligence*, 14, 165-203.
- Williams, M., Schofield, D., & Denby, B. (1998). The development of an intelligent haulage truck simulator for improving the safety of operation in surface mines. In J.-C. Heudin (Ed.), *Lecture Notes in Artificial Intelligence 1434 Proceedings of Virtual Worlds 98, First International Conference on Virtual Worlds* (pp. 337-344). Berlin: Springer.
- Wittig, T., & Onken, R. (1992). Knowledge based cockpit assistant for controlled airspace flight operation. In H. G. Stassen (Ed.), *Analysis, design and evaluation of man-machine systems 1992*. Oxford, UK: Pergamon.
- Woods, D. D., Patterson, E. S., Roth, E. M., & Christoffersen, K. (1999). Can we ever escape from data overload? A cognitive systems diagnosis. In *Proceedings of the Human Factors and Ergonomics Society 43rd Annual Meeting* (pp. 174-178). Santa Monica, CA: Human Factors and Ergonomics Society.
- Woods, D. D., Roth, E. M., & Pople, H. J. (1987). *Cognitive environment simulation: An artificial intelligence system for human performance assessment* (Nureg-CR-4862). Washington, DC: U.S. Nuclear Regulatory Commission.
- Wooldridge, M., & Jennings, N. R. (1995). Intelligent agents: Theory and practice. *Knowledge Engineering Review*, 10(2), 115-152.
- Wooldridge, M., Mueller, J., & Tambe, M. (Eds.). (1996). *Intelligent Agents Vol II (ATAL-95)*. Volume 1037 of Springer-Verlag's Lecture Notes in AI series. Berlin: Springer-Verlag.
- Wooldridge, M., & Rao, A. (Eds.). (1999). *Foundations of rational agency*. Dordrecht, The Netherlands: Kluwer.
- Wray, R. (2001). Soar-based technology: A perspective. In *Proceedings of 21st Soar Workshop* (pp. 45-47). Ann Arbor, MI: The University of Michigan Soar Group. ai.eecs.umich.edu/soar/workshop21/talks/author-index.html
- Wright, I. P., Sloman, A., & Beaudoin, L. P. (1996). Towards a design-based analysis of emotional episodes. *Philosophy Psychiatry and Psychology*, 3(2), 101-126.

- Yost, G. R. (1992). *TAQL: A Problem Space Tool for Expert System Development*. Unpublished doctoral dissertation, Carnegie-Mellon University, Pittsburgh.
- Yost, G. R. (1993). Acquiring knowledge in Soar. *IEEE Expert*, 8(3), 26-34.
- Yost, G. R., & Newell, A. (1989). A problem space approach to expert system specification. In *Eleventh International Joint Conference on Artificial Intelligence* (pp. 621-627). Menlo Park, CA: AAAI Press.
- Young, R. M. (1999). A zoo of browsable, runnable cognitive models. In D. Peterson, R. J. Stevenson, & R. M. Young (Eds.), *Proceedings of the AISB '99 Workshop on Issues in Teaching Cognitive Science to Undergraduates*. 25. Brighton, UK: The Society for the Study of Artificial Intelligence and Simulation of Behaviour.
- Young, R. M., & Barnard, P. J. (1987). The use of scenarios in human-computer interaction research: Turbocharging the tortoise of cumulative science. In J. M. Carroll & P. Tanner (Eds.), *Human Factors in Computing Systems and Graphics Interfaces (CHI & GI 87)* (pp. 291-296). New York: ACM.
- Young, R. M., Barnard, P. J., Simon, A. J., & Whittington, J. E. (1989a). How would your favourite user model cope with these scenarios? *ACM SIGCHI Bulletin*, 20(4), 51-55.
- Young, R. M., Green, T. R. G., & Simon, T. (1989b). Programmable user models for predictive evaluation of interface designs. In *Proceedings of CHI'89 Conference on Human Factors in Computing Systems* (pp. 15-19). New York: ACM.
- Young, R. M., & Lewis, R. L. (1999). The Soar cognitive architecture and human working memory. In A. Miyake & P. Shah (Eds.), *Models of working memory: Mechanisms of active maintenance and executive control* (pp. 224-256). New York: Cambridge University Press.
- Zettlemoyer, L. S., & St. Amant, R. (1999). A visual medium for programmatic control of interactive applications. In *CHI '99, Human Factors in Computer Systems* (pp. 199-206). New York: ACM

Author Index

Aasman	21, 62, 83, 87, 88, 89, 94, 95
Agre	10, 13, 83
Akkermans	98
Akyürek	83
Albus	33, 83
Alechina	31, 83, 94
Amalberti	41, 83
Anderson	4, 5, 12, 13, 16, 23, 24, 25, 35, 73, 77, 83, 85, 93, 99
Angros	71, 83
Anjewierden	98
Anzai	10, 83
Arnold	28, 93
Assanie	71, 83
Baddeley	11, 46, 84
Bainbridge	88
Barnard	28, 65, 68, 101, 102
Bartl	xiii, 33, 34, 84
Bass	16, 21, 35, 84
Bauman	70, 91
Baxter	16, 20, 25, 35, 64, 68, 84, 94, 97, 99
Baylor	17, 84
Beaudoin	102
Belavkin	13, 34, 55, 56, 75, 84
Bellotti	92
Bibby	10, 21, 28, 97
Bigus	66, 84
Blake	27, 99
Bloedorn	73, 84
Boff	6, 9, 11, 12, 27, 55, 84
Bovair	12, 85
Brazier	36, 85
Breuker	65, 101
Brooks	19, 21, 39, 73, 85, 90
Bunzon	15, 94

Burke	29, 85
Burton	27, 99
Busetta	36, 66, 85
Butler	70, 85
Byrne	12, 21, 22, 35, 44, 62, 78, 85
Cacciabue	15, 41, 42, 85, 90
Campbell	30, 85
Carbonell	96
Carpenter	11, 12, 87, 92
Ceranowicz	2, 25, 86
Cesto	85
Chabay	94
Chapman	13, 83
Charness	16, 17, 86
Chase	10, 16, 17, 32, 86, 95, 99
Cheng	21, 93
Chipman	10, 27, 35, 55, 66, 86, 88, 98
Chong	21, 22, 27, 35, 44, 55, 78, 85, 86
Christoffersen	18, 101
Cooper	34, 86
Corbett	73, 83
Corker	41, 93
Corlett	99
Cottam	70, 86
Coulter	22, 61, 92, 93
Crandall	10, 90
Crerar	13, 95
Crow	36, 86
Crowder	11, 86
Daily	12, 86, 94
Damasio	38, 56, 87, 99
Daneman	12, 87
Dautenhahn	100
Dawes	6, 11, 87
de Groot	16, 19, 20, 21, 32, 87
de Hoog	98
de Keyser	41, 87
Deblon	41, 83

Decortis	15, 85
Delaney	59, 87
Denby	73, 90, 98, 101
Derry	101
Detje	34, 87
Dörner	13, 33, 34, 84
Downes-Martin	73, 84
Drozdownicz	15, 85
Duchi	54, 93
Dunin-Keplicz	36, 85
Elkind	xiii, 5, 87
Elliman	13, 21, 29, 75, 84, 85, 87, 90
Elman	33, 87
Erds	98
Ericsson	10, 12, 17, 27, 28, 68, 86, 87, 96
Ernst	59, 95
Etzioni	36, 87
Farrell	16, 83
Fellous	34, 91
Feltovich	99
Feurzeig	73, 97
Fisher	28, 98
Franklin	36, 88
Freed	15, 41, 42, 44, 85, 88, 91
Frese	15, 88
Frijda	38, 95
Fürnkranz	89
Gaines	86
Gardin	95
Garlick	71, 100
Geddes	72, 88
Gernsbacher	101
Gigley	35, 55, 66, 88
Gilhooly	98
Glover	29, 88
Gluck	xiii, 25, 66, 88
Gobet	10, 16, 17, 19, 21, 32, 59, 69, 75, 87, 89, 93, 96

Goldberg	28, 89
Goss	xiii, 1, 94
Graesser	36, 88
Gratch	13, 44, 55, 70, 89
Gray	44, 85
Green	55, 58, 89, 102
Greenberg	72, 89
Grimes	21, 90
Hart	30, 90
Hatano	95
Heise	71, 90
Herschberg	89
Hill	21, 62, 90
Hirst	25, 70, 92
Hochberg	xiii, 5, 87
Hodgson	36, 85
Hoffman	10, 17, 18, 27, 90, 99
Holland	44, 90
Hollands	xiii, 73, 90
Hollnagel	41, 42, 43, 57, 90
Holyoak	44, 90
Howden	xiii, 36, 85
Howes	9, 16, 24, 75, 91, 97
Hudlicka	34, 55, 91
Huey	xiii, 5, 87
Huffman	71, 91
Hull	30, 85
Ioerger	54, 99
Jackson	30, 85
Jacobs	36, 91
Jansen	xiii, 17, 19, 21, 89, 91
Jaulin	95
Jeffries	92
Jennings	36, 88, 101
John	10, 23, 40, 41, 42, 91, 95
Jones	13, 20, 21, 22, 25, 27, 34, 35, 44, 61, 64, 70, 72, 77, 79, 91, 92, 93, 96, 97, 100
Josephson	54, 86

Just	11, 12, 68, 92
Kalus	25, 70, 92
Kanerva	33, 92
Kaufman	6, 83, 84, 85, 93, 100
Keane	98
Keirse	62, 99
Kenny	22, 61, 92, 93
Kieras	10, 35, 42, 68, 91, 92
Kintsch	10, 12, 17, 35, 87, 92
Kirsh	85
Kirsner	10, 40, 95
Kitajima	35, 92
Klein	10, 17, 92
Koenig	19, 20, 61, 92
Koriat	33, 92
Koss	22, 61, 62, 72, 92, 93, 99, 100
Kosslyn	19, 20, 61, 92
Krems	97
Kubat	89
Kuk	28, 93
Laguna	29, 88
Laird	12, 16, 22, 54, 61, 70, 71, 72, 77, 83, 86, 91, 92, 93, 94, 100
Lajoie	15, 94
Lane	21, 93
Langley	71, 80, 93, 100
Larkin	10, 25, 28, 60, 68, 70, 93, 94, 97
Laughery	41, 93
Lebiere	5, 12, 13, 16, 23, 25, 77, 78, 83, 85, 93
Lehman	23, 95
Leonardo	61, 100
Lesgold	15, 92, 94
Lewis	12, 16, 23, 66, 71, 90, 94, 102
Lieblich	33, 92
Lincoln	6, 9, 11, 12, 27, 55, 84, 91
Ling	68, 100
Logan	xiii, 29, 31, 38, 40, 43, 48, 75, 83, 94, 99, 100

Logie	98
Lonsdale	xiii, 62, 75, 94
Loral	3, 94
Lovett	12, 86, 94
Lucas	xiii, 1, 36, 85, 94
Mach	15, 94
MacKinlay	92
Major	41, 72, 97, 116
Marsella	13, 89
Massey	97
Masson	15, 85
Matessa	13, 83, 91
Mavor	xiii, 1, 4, 5, 6, 9, 12, 20, 21, 25, 26, 27, 31, 32, 35, 41, 43, 53, 59, 61, 62, 66, 96
McBride	13, 88
McClelland	28, 94
McNeese	xiii, 41, 98
Merrill	35, 100
Meyer	35, 92
Meyrowitz	10, 27, 86
Michalski	96
Michon	21, 83
Miller, C. S.	16, 94
Miller, G. A.	11, 94
Misker	62, 95
Mitchell	96
Miyake	12, 95, 102
Moffat	38, 95
Monk	13, 95
Mueller	39, 101
Müller	88
Musen	86
Mutter	97
Neisser	19, 21, 42, 61, 95
Nelson	23, 95
Nerb	59, 95
Newell	4, 5, 11, 12, 13, 16, 17, 28, 35, 40, 44, 68, 69, 70, 77, 78, 91, 93, 95, 102

Nhouyvanisvong	33, 98
Nielsen	10, 22, 40, 61, 92, 93, 95
Nilsson	30, 90
Nisbett	44, 90
Nordvik	15, 85
Nwana	35, 95
Ohlsson	10, 71, 93, 96
Okada	95
Ong	64, 96
Onken	72, 101
Patterson	18, 84, 101
Payne	24, 96
Pearl	30, 96
Peterson	35, 98, 100, 102
Pew	xiii, 1, 4, 5, 6, 9, 12, 20, 21, 25, 26, 27, 31, 32, 35, 41, 43, 53, 59, 61, 62, 66, 88, 96
Piaget	69, 96
Picard	13, 34, 56, 96, 99
Picton	21, 90
Pitrat	17, 96
Polson	35, 92
Pople	41, 101
Popper	4, 96
Pсотка	97
Pylyshyn	20, 96
Quinlan	71, 96
Quintanilla	88
Rao	39, 101
Raphael	30, 90
Rasmussen	14, 33, 41, 96
Rauterberg	18, 96
Ray	56, 100
Reason	41, 96
Reder	12, 16, 33, 59, 86, 87, 93, 94, 98
Remington	15, 41, 88, 91
Richards	33, 98
Richman	19, 32, 59, 61, 89, 96, 97

Riedl	19, 21, 62, 100
Rieman	9, 23, 97
Ritter	xiii, 10, 13, 16, 20, 21, 22, 23, 25, 28, 34, 35, 44, 55, 59, 60, 61, 62, 64, 68, 69, 70, 72, 73, 75, 77, 78, 80, 84, 85, 87, 89, 91, 92, 93, 94, 96, 97, 98, 100
Rönnquist	xiii, 36, 85
Root	30, 85
Rosenbloom	12, 13, 72, 77, 93, 98, 100
Roth	18, 41, 101
Rouse	68, 98
Roytam	70, 98
Rumelhart	28, 94
Saariluoma	16, 98
Salvendy	93
Salvucci	21, 22, 98
Sandberg	28, 68, 101
Sanderson	28, 41, 98
Sasse	13, 95
Sauers	16, 83
Scheftic	94
Schmid	97
Schobbes	85
Schofield	73, 98, 101
Schraagen	27, 98
Schreiber	11, 65, 98, 101
Schunn	33, 98
Schwamb	62, 67, 72, 99, 100
Seif El-Nasr	54, 99
Sekuler	27, 99
Shadbolt	10, 18, 27, 36, 70, 75, 86, 90, 95, 98, 99
Shafto	15, 80, 88, 92
Shah	12, 95, 102
Shalin	27, 98
Shea	36, 91
Sheldon	13, 88
Shrager	10, 83

Siegler	59, 99
Singley	77, 99
Skidmore	72, 89
Skinner	69, 99
Sloman	xiii, 13, 31, 37, 38, 39, 40, 60, 75, 94, 96, 99, 100, 102
Small	72, 89
Smith	86
Soto	35, 92
Spada	59, 95
Spector	28, 66, 100
St. Amant	19, 21, 62, 100, 102
Stassen	101
Staszewski	19, 32, 59, 87, 89, 96, 97
Stern	56, 100
Stevenson	98, 102
Stroffolino	33, 98
Taatgen	62, 87, 88, 89, 94, 95
Tabachneck-Schijf	61, 100
Tambe	39, 72, 100, 101
Tanabe	95
Tauber	92
Tenney	28, 66, 100
Thagard	44, 90
Thomas	6, 84
Treur	36, 85
Uiterwijk	89
van den Herik	89
van Lent	71, 100
VanLehn	71, 100
Vera	40, 91
Verbrugge	36, 85
Vicente	16, 30, 101
Wallace	64, 101
Wallach	59, 98
Wang	35, 101
Weare	29, 85
Weld	36, 87

Whalen	14, 101
Whittington	66, 102
Wickens	11, 101
Wielinga	65, 98, 101
Wilkins	17, 101
Wilson	99
Winne	96
Wittig	72, 101
Wood	21, 34, 91
Woods	18, 41, 87, 101
Wooldridge	36, 39, 88, 101
Wray	23, 102
Wright	48, 84, 102
Wysotzki	97
Yen	54, 99
Yerkes	55
Yost	65, 68, 70, 102
Young	9, 12, 16, 20, 23, 55, 62, 64, 65, 66, 75, 84, 91, 92, 97, 98, 100, 102
Zaff	41, 98
Zenyah	72, 89
Zettlemoyer	19, 102
Zhang	35, 101

Subject Index

Major terms are flush left. Minor or subterms are indented, duplicate terms in italics.

A* algorithm	29, 30, 31, 48, 81
A*epsilon	30
A* with Bounded Costs (ABC)	31, 81
Able	60, 70
ACM	71
active perception	21
Adaptive Control of Thought – Rational (ACT-R)	4, 12, 16, 22, 23, 27, 32, 34, 35, 38, 40, 44, 55, 59, 60, 61, 62, 65, 67, 77- 81
Adaptive Control of Thought - Rational with Perceptual-Motor component (ACT-R/PM)	22, 35, 38, 62, 78, 81
adversarial problem solving	15-17
affordances	37
aggregation of behavior	9, 26
AIDE	4
air traffic control	21, 41, 68
Amulet	62
amygdala	14
APEX	41-42, 45, 47, 49, 51, 81
attention	5, 12, 14, 16, 22, 38, 39, 61
augment perception	20
automated operator support	72
automatic model generation	71, 73
back-propagation algorithm	32
backward-chaining reasoning	10, 60
behavior moderators	5,9, 12-14, 31, 37-40, 55, 56, 73
Belief-Desire-Intention (BDI) agents	36, 45-52, 81
C	51
C++	46, 68
CAMERA	61
CAPS	12
CASSY	72
categorization	19, 37, 38, 39
cerebellum	33, 45
CES	41, 81
chains of actions	58
chess	16, 17, 32, 52
chess experts	32
CHIRP	58, 81
CHREST	19, 32, 81

chunk	16, 19, 21, 22, 23, 31, 32, 48, 52, 60, 79, 81
Cirrus	71, 100
civilians, non-combatants, and white forces	1, 18
COGENT	1, 34, 44, 45-52, 70, 77
cognitive architecture	iv, xiii, 4-6, 13, 22, 23, 26, 27, 29, 31, 33, 35, 37, 43, 44, 55, 56, 59, 61, 63-67, 73, 77, 80
cognitive engineering	21, 41-43
Cognitive Environment Simulation (CES)	41, 81
cognitive models	1, 2, 4, 12, 21, 23, 25, 31, 33, 34, 42, 60, 62, 73
Cognitive Reliability and Error Analysis Method (CREAM)	41, 43, 57, 58, 81
Cognitive Simulation Model (COSIMO)	41, 81
comparisons with data	52-54, 66, 68-70
concept formation	19, 31, 32, 59
Confidential Human Factors Incident Reporting Program (CHIRP)	58, 81
connectionism	9, 32, 39, 47, 55
consciousness	14, 38
constraints on cognitive models	3, 16, 20, 24, 30, 31, 48, 77
Contextual Control Model (CoCoM)	42, 43, 81
creating models	5, 9, 22-24, 27, 28, 51, 53, 65, 70, 71, 77
Critical Decision Method	10
data gathering for testing models	27-28, 53
decision making	5, 6, 10, 17, 22, 42, 48, 60
decision trees	71
declarative memory	47, 78, 79
Defence Evaluation and Research Agency (DERA)	xiii, 10, 40, 75, 81
Defence Science and Technology Organisation (DSTO)	xiii
Delphi	33, 51, 62
Diligent	71
DIS protocol	2, 3, 24, 63, 81
discrimination net	31, 32, 45, 46, 47, 51
display-based skill	24
distractions	12
distributed artificial intelligence	35, 36, 50, 67
doctrine	2, 10, 14, 15, 18, 33, 35, 53, 56, 63
non-doctrinal performance	13, 14
documenting models	40, 44, 65
domain knowledge	29, 43, 53
domain ontologies	36

- embedded assistants 72
- emotion iv, 5, 12-14, 20, 26, 31, 33-34, 37-40, 44, 45, 54, 55-56, 59, 60, 73, 84, 91, 95, 97
- engineering models 41-43, 45-52
- engineering psychology 41, 68
- entertainment 6, 73, 86
- EPAM 16, 19, 21, 31-32, 44, 45-52, 55, 59, 61, 4, 22, 44, 81
- EPIC
 - and ACT-R 35
 - and Soar 27, 35
- ergonomics vii, 20
- errors
 - erroneous action 56, 57, 58
 - erroneous behavior 9, 16, 56, 57, 59
 - error 10, 11, 15, 16, 18, 33, 37, 42, 48, 56, 57, 58, 66
 - error rates 10
 - how to manage error 15
 - human error 42, 57
 - mistakes 13, 15, 38
 - models of 16, 56-57
 - taxonomy of human error 57
 - training about 15
- evaluation function 17, 28, 29-31
- expertise 10-11, 12, 16, 18, 20, 36, 60
- experts 15, 17, 18, 20, 23, 27, 32, 35, 36, 52, 58, 59, 60, 73
- exploratory learning 23
- exploratory search 24
- eye movements 19, 22, 45, 46, 51
- fatigue 5, 13, 14, 55, 56
- fear 14, 56
- feature extraction 20
- feeling of knowing 33
- flagship task 63
- framework, theoretical 4, 5, 14, 24, 30, 36, 38, 39, 41, 45, 63, 65, 67, 77
- Garnet 62
- generic tasks 65
- Genetic Algorithms (GAs) 28-29, 43, 69, 70
- graphic interface 34, 39, 50, 51, 62
- graphic representation of models 65
- Hazard Monitor 72
- heuristic search 29-31, 77

heuristics	17, 29
high-level modeling languages	65, 70
Higher Level Architecture (HLA)	67
higher-level vision	20, 21, 61
hill-climbing	29, 48
human factors	20, 42
Human Systems Information Analysis Center (HSIAC)	iv, vii, xiii
hybrid architectures	iv, 1, 4, 35, 44, 55, 66
ID3 learning algorithm	71
IDXL	23
incident reports	42, 58, 81
Individual Data Modeling (IDM)	69
individual differences	12, 13, 17, 20, 59, 69
information overload	18
Infosleuth architecture	36
Instructo-Soar	71
insubordination	13
integrated architectures	23, 62
integrating model components	22-23, 65
integrating models with simulations	iv, xiii, 1, 2, 22, 61-64, 72, 73
integrating psychology theories	22
intelligent architectures	41
JAVA	36, 45, 46, 47, 50, 51, 62, 66, 67, 68, 85
JAVA Agent Construction Kit (JACK)	1, 36, 45-52, 66
knowledge acquisition	18, 27, 28, 36, 70
Knowledge-Based Systems (KBSs)	35, 36, 65, 72
knowledge level	25, 35, 36, 70
learning	9-11, 21, 22, 23, 24, 31, 32, 33, 37, 45, 46, 48, 49, 54, 55, 59, 60, 61, 71, 77, 79
learning to model	23, 68-69, 72
LICAI	35
limited working memory	12, 16, 17, 24
Lisp	40, 50, 51, 60, 62, 70
Long-Term Memory (LTM)	12, 17, 29, 31, 46, 52, 61
long-term working memory	12, 82
machine learning	55, 70, 71, 72
meta-management of cognition	40, 48
model, cognitive	
comparing models (bake-offs, matrix exercises)	65-66
of emotions	12-13, 17, 33-34, 37-40, 55
methodology of	68-71
model accuracy	18, 26, 54

model implementation language	51
model libraries	23-24, 34, 40, 51, 65, 70
modeling interfaces	25, 70, 71-72
of personality	59
of teamwork	1, 5, 50, 59
of urban terrorism	9
Modular Semi-Automated Forces (ModSAF)	2-4, 24, 62-64, 72, 82
motivation	9, 12, 13, 20, 33, 45, 55, 56
motor behavior	19, 20, 21, 22, 33, 35, 37, 46-49, 52, 78
models of	21, 35, 46, 47-49, 52, 78
Multi-Layer Perceptron (MLP)	32, 60
multicriteria heuristic search	29-30
multiple strategies	17-18, 59
multitasking	5, 9, 41, 42, 49, 52
NASA	42, 66, 81
NASA Test Director	23
natural language	23
Naturalistic Decision Making (NDM)	10, 60
network,	
computer	3, 63, 64, 67
neural	4, 32-33, 37, 46, 47, 61
noradrenalin	14
Nottingham Functional Interaction Architecture	22, 35, 44, 62
novices	10, 11, 15, 18
objective function (see evaluation function)	
Observe-Soar	71
Office of Naval Research (ONR)	xiii, 35, 55, 66
operant conditioning	48
ore mining, application of models	73
panic	18
pattern recognition	16, 17, 18, 32, 44, 45
PDP	23, 46
PDP++	23
perception	6, 9, 10, 14, 17, 19-22, 36, 40, 45, 62
perception models	18, 19-22, 35, 40, 42, 45, 60, 61-62, 64, 72, 78
and problem solving	20-21
perceptual cycle	20, 21, 42, 61
personality	13, 17-18, 38, 55, 59-60, 73
phonological loop	11
physical models	3, 39, 63, 64
physiological influences on cognition	14, 33, 56, 57
planning	5, 13, 14, 17, 21, 23, 28-31, 36, 38-39, 42-43, 44, 48, 58, 60-61, 70
Pop-11	40, 46, 47, 51

Poplog system	40
power law of learning	77
practice	10, 12, 15, 32, 37
problem solving	10, 13, 15, 16-17, 19-24, 29, 30, 31, 35-36, 44, 48, 55-56, 59-60, 77-79, 84
problem-solving styles	10, 11, 13, 17-18, 38, 59
procedural knowledge	47, 77-79
production rules (productions)	22, 46, 47, 48, 61, 78, 79
production systems	4, 21, 34, 56, 68, 81
progressive deepening	17
Prolog	40, 47, 51
protocol analysis	28, 68
protocols	27, 71
verbal protocol analysis	28, 68
PSI	22, 33, 34, 44, 45-52, 60, 77
Purple Link	2
RAF Institute of Aviation Medicine	58
Rapid Decision Making (RDM)	10, 17, 60
models of	60, 70
reactive architectures	13, 19, 21, 37-40, 60, 78
reactive behavior	13, 14, 19, 23, 33, 37, 38, 39, 40, 55, 56, 60, 78
recognition-based problem solving	24, 60, 70
recurrent nets	33
retina	20
reuse of cognitive models	1, 22-24, 34-36, 44, 62-63, 65, 70
review of simulation systems	62-63
robotics	19, 21, 22, 72
route planning	30-31
rules (also see productions)	23, 46, 47, 49, 56, 59, 60, 71, 78, 79, 96
satisficing	16, 17, 30
schemas	21-22, 35, 39, 45, 47, 48, 61
search	10, 11, 16-17, 23-24, 28-31, 33, 60, 77, 46
semantic net	46
sensitivity analysis	15, 18, 25, 73
sequences	3, 16, 24, 33, 46, 47, 49, 58, 60
Short-Term Memory (STM)	11, 16, 17, 38, 46, 49, 78
Sim_Agent Toolkit	37-40, 44, 45-52, 55, 64
Simplified Model of Cognition (SMoC)	42-43
situation assessment	13, 18, 48
situation awareness	5, 20, 48, 60

- Soar 4, 12, 16, 21, 22, 23, 26, 27, 29, 32, 35, 38, 40, 44, 53, 54, 55, 59, 60, 61, 62, 64, 65, 67, 70, 71, 72, 77-80
- Soar Quake-bot 54
- Sparse Distributed Memory (SDM) 33, 45-52, 60
- speaking 23
- speed of mental processing 10, 16, 39
- speed of movement 3, 54
- Standard ML 51
- STOW '97 2
- strategies 9-11, 16-18, 24, 28-29, 38, 41, 59, 73
- stress 10, 12, 14, 20, 55-56, 59, 60
- summary measures 33-34, 53, 66
- synthetic environments iv, xiii, 1-3, 5-7, 9, 11, 15, 16, 17, 18, 20, 22, 24, 26, 28-29, 35, 53-55, 58, 60-64, 67-68, 71, 72-73
- synthetic forces 1-3, 6, 10, 13, 17, 18, 25, 63
- System Response Generator (SRG) 42
- Tabu search 29, 60
- Tac-Air system 64, 72
- tactics 2, 3, 9, 10, 43, 55
- Tank-Soar 64
- task simulations 2, 53, 62, 63-64, 66-68
- Tcl/Tk 62
- techniques for building and using models iv, 25, 65, 68-71
- tertiary emotions 38
- thinking by analogy 33
- time pressures 10, 14, 17-18, 41
- tip-of-the-tongue 33
- training 1-3, 9-11, 15, 18, 23, 25, 26, 54, 58, 61, 63, 69, 73
- transfer of knowledge 10, 37, 54
- Turing test 54
- tutorials on modeling architectures 34, 44, 68, 70
- Unified Theory of Cognition (UTC) 4, 13, 22-23, 69
- unified theory of emotion 13, 55
- usability of models iv, xiii, 1, 4, 6, 25, 41, 44, 53, 65, 67, 68-73
- user interface 25, 31, 51, 64, 65, 68, 70, 72
- validation of models 26, 52-54, 65, 66, 68-69
- variation in behavior 13, 17, 20, 54, 55, 59, 62
- virtual reality, models used in 73
- Visual Basic 62
- visual scanning 21, 23, 46
- visual-spatial scratchpad 11
- Windows 50, 62

working memory	9, 11-12, 16, 24, 46, 56, 62
X windows	50
Yerkes-Dodson law	55